NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

AN EVOLVING ASYMMETRIC GAME FOR MODELING INTERDICTOR-SMUGGLER PROBLEMS

by

Richard J. Allain

June 2016

Thesis Advisor: David L. Alderson
Second Reader: W. Matthew Carlyle

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			*Form Approved OMB No. 0704–0188*

1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** June 2016	**3. REPORT TYPE AND DATES COVERED** Master's thesis

4. TITLE AND SUBTITLE AN EVOLVING ASYMMETRIC GAME FOR MODELING INTERDICTOR-SMUGGLER PROBLEMS	**5. FUNDING NUMBERS**
6. AUTHOR(S) Richard J. Allain	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000	**8. PERFORMING ORGANIZATION REPORT NUMBER**
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A	**10. SPONSORING / MONITORING AGENCY REPORT NUMBER**

11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol number ____N/A____.

12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited	**12b. DISTRIBUTION CODE** A

13. ABSTRACT (maximum 200 words)

We propose a novel network interdiction model that reconciles many operational realities identified by military literature. Specifically, we conduct network interdiction within a dynamic network under partial information, using incomplete feedback and allowing two-sided adaptive play. Combining these aspects in an evolving game, we use optimization, simulation, and stochastic models to achieve a hybrid model. Modeling some currently underrepresented martial problems in this way makes it possible to highlight otherwise obscure relationships between policy and outcome, and to discover emergent effects, such as deterrence. As an example of this class of problems, we consider the struggle between a smuggler and interdictor. The smuggler seeks to maximize the amount of forces and materiel infiltrated from an origin to destination. The interdictor seeks to minimize this smuggler flow. Using two simple examples of an illicit-trafficking network, we demonstrate how to use these quantitative models within such an interdictor-smuggler context to (1) evaluate the value of seizures as a proxy for smuggled materiel, (2) assess the value of exploration, and (3) provide decision makers with practical ways to better allocate resources and increase effectiveness.

14. SUBJECT TERMS dynamic network interdiction, smuggling, evolutionary games, incomplete and asymmetric information, online learning, dynamic stochastic programming, adaptive risk management	**15. NUMBER OF PAGES** 141
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	**18. SECURITY CLASSIFICATION OF THIS PAGE**	**19. SECURITY CLASSIFICATION OF ABSTRACT**	**20. LIMITATION OF ABSTRACT**
Unclassified	Unclassified	Unclassified	UU

NSN 7540–01-280-5500 Standard Form 298 (Rev. 2–89) Prescribed by ANSI Std. 239–18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

AN EVOLVING ASYMMETRIC GAME FOR MODELING INTERDICTOR-SMUGGLER PROBLEMS

Richard J. Allain
Major, United States Marine Corps
B.S., University of Florida, 1999

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

NAVAL POSTGRADUATE SCHOOL
June 2016

Approved by: David L. Alderson
Thesis Advisor

W. Matthew Carlyle
Second Reader

Patricia A. Jacobs
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

We propose a novel network interdiction model that reconciles many operational realities identified by military literature. Specifically, we conduct network interdiction within a dynamic network under partial information, using incomplete feedback and allowing two-sided adaptive play. Combining these aspects in an evolving game, we use optimization, simulation, and stochastic models to achieve a hybrid model. Modeling some currently underrepresented martial problems in this way makes it possible to highlight otherwise obscure relationships between policy and outcome, and to discover emergent effects, such as deterrence. As an example of this class of problems, we consider the struggle between a smuggler and interdictor. The smuggler seeks to maximize the amount of forces and materiel infiltrated from an origin to destination. The interdictor seeks to minimize this smuggler flow. Using two simple examples of an illicit-trafficking network, we demonstrate how to use these quantitative models within such an interdictor-smuggler context to (1) evaluate the value of seizures as a proxy for smuggled materiel, (2) assess the value of exploration, and (3) provide decision makers with practical ways to better allocate resources and increase effectiveness.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF FIGURES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

DFS	Depth First Search
DOD	United States Department of Defense
DOE	Design of Experiments
GAMS	General Algebraic Modeling System
GB	Gigabyte
GHz	Gigahertz
HB	Hostile Base
I-I/S	Interdictor—Interdictor/Smuggler
ILP	Integer Linear Program
ISIL	Islamic State of Iraq and the Levant
OLH	Orthogonal Latin Hypercube
PC	Personal Computer
RAM	Random Access Memory
RAND	Project Research ANd Development
RS	Reverse Star

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

With the rise of war by proxy and growing transnational terrorism and organized crime; smuggling and infiltration, and efforts to disrupt these activities through interdiction have become central to many nations' security. Failure by a government to guard against both infiltration and smuggling leads to direct threats to these governments' sovereignty. To help tactical decision makers better address these threats, we consider a resource allocation problem faced by the interdictor, specifically, where to position limited combat power along suspected infiltration routes within the operating area over time.

The interdictor must make these decisions using incomplete information on the location of in-transit smuggler goods and the cost and capacity of the smuggler routes. Further compounding the problem, we assume the smuggler is both intelligent and adaptive. The interdictor's objective is to best-disrupt the smuggler's lines of communication by minimizing the amount of forces and materiel travelling from the smuggler's safe haven to a destination. Opposing this aim, the smuggler's goal is to maximize the flow of forces and materiel to the destination.

Previous research into martial contests, such as this interdictor-smuggler context, has made various strong assumptions in the name of tractability. These persistent assumptions include perfect information, an unchanging environment, and non-adaptability. We believe these assumptions do not adequately represent problems within this interdictor-smuggler context and many other tactical scenarios.

Our approach is bottom up. We merge strong aspects of traditional game theory, optimization, stochastics, and simulation. First, we design a game with simple rules to examine patterns of interest and search for recognizable emergent behaviors. This game is instantiated in a terminating discrete event simulation. By repeating this simulation in a variety of configurations, each continued over a finite number of time steps, we study the time dynamics of our problem in a spectrum of situations. Through this evolution of play,

we examine the robustness of interdictor resource allocation policies and smuggling tactics. Lastly, we search for emergent relationships during play.

Our computational study begins with traditional relaxations that include perfect information for the interdictor and smuggler. We then demonstrate the effect of private information for each player. Private information requires the interdictor and smuggler to make estimations of their antagonist's current state and designs. These estimations combined with resource allocation decisions introduce adaptive play under uncertainty. We show that adaptive play under uncertainty causes a massive perturbation to both game play and outcome.

Next, we simulate games configured with several specific interdictor and smuggler policies using design of experiments. We consider a spectrum of interdictor policies that vary both the total resource budget and blend of resource types. The type of interdictor resource varies across a spectrum describing its visibility to the smuggler. The spectrum ranges from overt (highly visible) to covert (less visible). We also vary the scheme by which the smuggler allocates forces and materiel for infiltration across shipments. During analysis, we examine each game's time dynamics and end-results through the lens described above.

We demonstrate our model using two interdictor-smuggler problem instances, exposing a number of practical insights useful to decision makers selecting resource allocation policies within a counterinsurgency or counter-illicit trafficking setting. The first problem instance portrays a smuggling network with a large number of possible parallel routes, the second a network with fewer alternative paths but significant depth.

We find that *when allotted a small number of interdiction resources relative to the number of available smuggling routes, the interdictor should employ these resources in a manner highly visible to the smuggler*. Overt deployment will disrupt smuggled forces and materiel primarily through deterrence. Using our model, we are able to assess the value of deterrence, given other policy options.

As the interdictor's resource budget increases, deploying forces less visible to the smuggler alongside those highly visible to the smuggler becomes more effective. These

highly visible forces should then be used to herd smugglers into waiting ambushes of less visible interdiction forces. In contrast to overt policies, pure covert policies primarily achieve their effect by actually seizing the smuggled flow. Deterrence is far less important under pure covert schemes. For these policies, information is key. Pure covert policies encourage the interdictor to target arcs deeper within the smuggling network more precisely.

Intuitively, *one might expect that the more the interdictor discovers the network, the more he disrupts smuggling, but we find that this is not the case*. Even so, if discovery of the physical structure of the network is important, interdictor policies with heavy allocations to less-visible forces are best suited to the task.

Finally, *we show that the amount of seized materiel is a poor proxy for the total amount of smuggled flow*. The relationship between these factors is very inconsistent. In policies where deterrence is high, the amount of seized materiel provides almost no information on the actual amount of unseen materiel successfully smuggled.

Beyond the direct findings listed above, we reach two deeper conclusions that have implications for research into martial contests, such as this interdictor-smuggler context:

(1) A range of realistic, complex behaviors can emerge from the interaction of two hostile, intelligent agents acting under simple rules within a dynamic environment of uncertainty and danger.

(2) We can gain insight into these complex situations through heuristics that combine complementary optimization, stochastic, and game-theoretic models under an umbrella of simulation.

This study is meant to be a prototype, demonstrating the power of a hybrid model. As a prototype, it is not without limitations. We make several assumptions on the method and speed by which hostile agents might evolve in a martial context. Additionally, we do not include the advantage obtained by interrogation and exploitation after the capture of enemy forces and materiel. Lastly, the duration of our computational cases is necessarily finite.

Future research could address these assumptions or admit real-world data to craft a wider array of smuggling networks or specific instances of interest. Under a broader set of configurations, models similar to ours could prove to be a great aid to training inter-agency decision makers and their staff. That training could encourage unique perspectives and seed important questions that might expose highly non-intuitive and indirect ways of influencing the outcome and assessing performance during real interdiction or counter-trafficking missions.

ACKNOWLEDGMENTS

What is to give light must endure burning.

—Victor Frankl
Man's Search for Meaning

Numerous mentors brought light forth from this thesis. Foremost among them is my advisor, Professor David L. Alderson. I owe a great deal of my education at the Naval Postgraduate School to him. He granted audience to my energetic, yet ill-focused ideas one year ago, and showed me how to blend my combat experience and academic rigor into a rational argument connected by a strong narrative—all written in active voice. My gratitude will be expressed by carrying forth the hard-won lessons and applying them with vigor.

I also want to thank Professor W. Matthew Carlyle for introducing me to the finer points of network interdiction models and game theory. He "taught me to fish," providing the tools to grow the model herein to its final form, and always found the time in his busy schedule to mentor me.

I also want to express my gratitude to Professors Samuel E. Buttrey, Robert A. Koyak, and Lyn R. Whitaker, the motley crew of statisticians within the Operations Research Department, who patiently entertained my myriad and sometimes repeated questions with poise. Without them, my analysis would have stalled well short of my goal. They helped me leverage simplicity alongside the most complex statistical models to yield insightful results.

In addition, I want to thank Professors Thomas W. Lucas and Paul J. Sanchez for suggesting a number of cutting-edge designs of experiments. The model within this thesis would have been intractable were it not for the deliberate techniques to which they introduced me. My gratitude also extends to Professor Ralucca Gera for the solid foundation in graph theory and network science that provided a wider perspective for this

thesis, and Professor Michael Atkinson for his support in researching Drug Trafficking Organizations.

Others who provided assistance on this journey include Captain Jeffrey Hyink and Lieutenant Colonel John Alt, who helped crunch these 141 pages into 8.5 minutes without loss; the coaches at the Graduate Writing Center who were instrumental in helping me find a clear voice in several sections of this thesis; and the professionals within the Thesis Processing Office who helped with the all-too-important details. Their dedication is under-valued; know that it is appreciated.

I also want to thank my classmates for teaching me a great many things both large and small. Their words of encouragement and humor in challenging times always lightened the load.

Lastly, I would not have been able to present this thesis without the devotion of my family. My lovely wife helped keep my priorities straight while creating time for this thesis in an equally busy schedule. My children showed me the powerful simplicity underlying many things but still managed to inspire a sense of awe. I would have neither seen nor taken this step without them.

I. INTRODUCTION

This thesis describes, develops, and exercises a mathematical model of a contest between a smuggler and an interdictor in an effort to study the allocation of interdiction resources and choice of smuggling tactics. We use this model to develop insights into different structures of such conflicts.

A. BACKGROUND

In the wake of two world wars, few modern nation states now choose to face one-another directly. The last 50 years have instead revealed a distinct rise in war by proxy. Malign non-state groups have also gained trans-national influence, threatening the strategic goals of many established nations. In response, these nations have increasingly turned elements of national power against such antagonists. Infiltration or smuggling and efforts to counter them are important features of these contests.

The U.S. DOD defines *infiltration* as the movement of small groups or individuals into a contested area by avoiding enemy contact (Marine Corps Combat Development Command 2001). Similarly, *smuggling* is defined as the "clandestine transportation of goods or persons past a point where prohibited…in violation of the law or other rules" (Lehman et al. 2004). Therefore, infiltration and smuggling are closely related. Each action is characterized by the contested movement of people, materiel, and information from a safe haven to a target operating area where these items find payoff. In armed conflict, the payoff is the focused application of force in time and space to create local advantage. The force might manifest as an improvised explosive device, kidnapping, ambush, or outright assault. In both war and peace, trafficked illicit goods are shepherded to areas of low supply and high demand. The payoff in such areas might be the sale of controlled narcotics, arrival of illegal migrants, or distribution of counterfeit currencies. The taxation of consumer goods smuggled from Dubai to Pakistan through Afghanistan provided an enormous portion of the Taliban's funding stream in the late 1990s and early 2000s (Rubin 2000). It is apparent that a government must guard against both smuggling and infiltration to maintain sovereignty.

In a failed or failing state, the government lacks sufficient influence to enforce control over some or all of its territory. Malign actors can further destabilize the weak government by fomenting an insurgency. Counterinsurgency operations are one method the host government can employ to resist the destabilizing agents. Counterinsurgency operations consist of offensive, defensive, and stability activities (United States Department of the Army 2008). Both the insurgent and counterinsurgent use unconventional warfare in pursuit of their respective aims.

Smuggling and infiltration are important features of unconventional warfare. Faced with a more numerous and better-equipped opponent, the insurgent's challenge is to create a temporary advantage through asymmetric means. The asymmetry involves a brief concentration of forces and materiel. Because the insurgent must maintain a low signature, maintaining a large, local stock of resources is risky. The insurgent must instead keep his primary sources of supply at a distance and disperse in the face of massed counterinsurgent forces. The high dispersion of both insurgent forces and materiel requires active and consistent lines of communication with a supporting organization for the insurgent to maintain any tactically significant level of activity.

These lines of communication usually connect a safe haven (*source*) and target operating environment (*target*). Within the source, sanctuary can be provided by political or military sponsors or even exceptionally difficult terrain. In either case, the source lies beyond the operational reach of the counterinsurgent. An intermediate area (a *network* of *routes*) connects the source and target. In the intermediate area, the insurgent can infiltrate forces and supplies. Of equal importance, intelligence and casualties can be transported away from the target or *exfiltrated*. Both the aforementioned infiltration and exfiltration are accomplished with the aid of a smuggler. For the smuggler, the intermediate area comprises his trade routes (Figure 1).

In the intermediate area, threat and opportunity meet. Here also, the counterinsurgent can act as interdictor, executing tactical action to disrupt the smuggler's activities, and thus the insurgent's lines of communication. Often of limited tactical interest, the intermediate area is strategically vital to both the insurgent and counterinsurgent. The success of interdiction or smuggling in the intermediate area can

isolate or relieve a contested area elsewhere, proving decisive. Because of the close relationship between smuggling and infiltration in the tactical context we describe, this thesis treats these terms as tacitly equivalent and we use the terms interchangeably.

Figure 1. The Source, Target, and Intermediate Area.

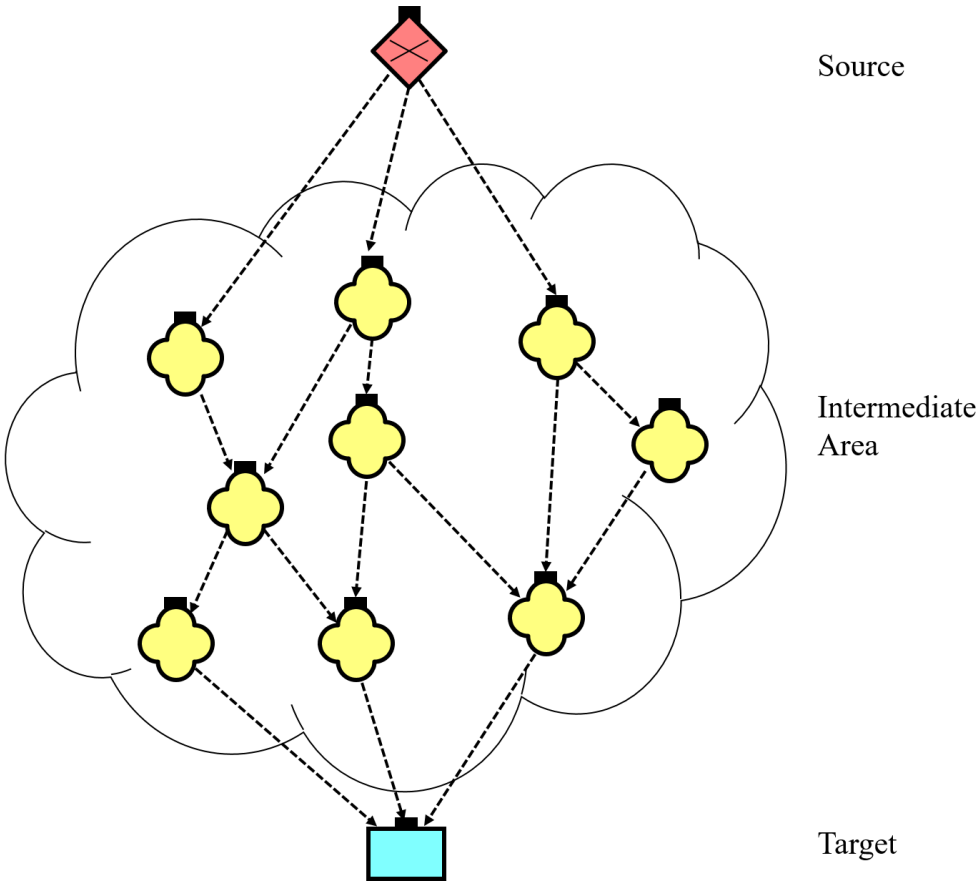

Lines of communication connect a safe haven (*source*) and target operating area (*target*) via a contested intermediate area (*network* of *routes*). Smugglers facilitate the movement of forces and materiel to forward insurgents through the network of routes.

Smuggling routes composing the insurgent's lines of communication are often small in number and change only after proving untenable. Operational experience ranging from the Vietnam War to more modern conflicts in the Middle East supports this point. The Ho Chi Minh trail was an essential lifeline for both the North Vietnamese Army and their proxies in the south, the Viet Cong (Prados 1999). Significant interdiction efforts by allied forces failed to alter this route despite practical alternatives (Prados 1999). In

Operation Iraqi Freedom, infiltration routes, or "ratlines," also tended to be highly consistent despite vast expanses of traversable desert. The constancy of insurgent lines of communication across the Durand Line, a line that separates Afghanistan and Pakistan, is similarly long-observed. Fighters, weapons, minerals, raw ores, and a vast spectrum of consumer goods still flow unabated over the same ground used to illicitly cross the Durand Line since the 1890s (Omrani 2009).

The above consistency makes lines of communication ripe for interdiction operations but poses a new issue. The sheer length and number of routes prevents would-be interdictors from maintaining persistent coverage of their entirety. Temporary outposts may be constructed, but smugglers quickly adjust routes around these obvious obstacles. Therefore, mobility is key and requires a large commitment of resources to even produce episodic concentrations of effectual combat power. The tactical problem is one highly sensitive to information and timing.

The smuggler and interdictor are thus locked in an asymmetric struggle characterized by extremely transient actions. A highly aggregate view of these actions might provoke a perspective that the interdictor-smuggler struggle is relatively stationary, each contest a repetition of the last with only some variance in the outcome. However, at the tactical level no combat is stationary; similar actions evoke wholly new responses when repeated. The new responses then generate new circumstances under which the interdictor and smuggler meet. A highly aggregate view that discounts the path-dependence of combat is seldom informative to problems of directing tactical action.

Tactical actions are focused and do not occur with consistent tempo. Preparation for some interdiction missions requires days, weeks, or even months. In execution, the resulting engagement between smuggler and interdictor is often resolved in seconds or minutes. The statistical average level of combat activity does not exist in the above reality. Each interdiction and smuggled shipment is unique but related to those that came before it. Both the interdictor and smuggler face a series of linked tactical problems, not one problem repeated in time.

Interdictors can employ combat power in various forms to address these tactical problems. Random patrolling, outposts, aerial reconnaissance, harassing fires, and ambushes are past examples of combat power employed in support of interdiction. No matter the specific form, these actions lie on a spectrum from overt to covert. The level of observability depends on the interdiction effort. A patrol conducted in the desert by armored vehicle would be highly overt. Dust clouds of approaching vehicles are visible for tens of miles. Conversely, high-altitude reconnaissance assets would be nearly invisible to a smuggler, and thus covert. Some activities, such as small ambush patrols, might lie somewhere in between these extremes.

Smugglers have their own spectrum of methods to maintain freedom of maneuver on their critical lifelines. They often use camouflage and concealment in non-traditional ways by disguising their activities as the movement of licit goods. Sophisticated observation networks warn smugglers of impending threats and offer direction to safer passage. The size of smuggled shipments also varies greatly, ranging from entire truckloads of weapons to information passed on a single flash memory card.

Both the interdictor and smuggler confront a situation full of complexity and ambiguity. A difference between orientation and reality is assured. The individual ability of the interdictor and smuggler to reconcile the difference between perception and reality by adapting assumptions and evaluating information within an evolving context is a significant factor in the conduct and result of the interdictor-smuggler struggle. Any analysis of an interdictor-smuggler problem must then consider the scenario's history as context. If adjusting perception to reality is paramount to success in the interdictor-smuggler context, it also suggests that the uncertainty level of the interdictor or smuggler may only have meaning in view of the other's relative information level.

Recent events show the importance of this interdictor-smuggler problem in stark relief. Within the last two years, criminal networks aided the infiltration of over 36,000 foreign fighters into Syria and Iraq (see Figure 2), swelling the Islamic State's ranks (Dilanian 2016). As of 17 August 2015, over 250,000 people died and 13,500,000 fled as a direct result of the fighting, creating a global humanitarian crisis (United Nations Officer for the Coordination of Humanitarian Affairs 2016). The impact of infiltration

was also apparent in Paris (2015) and Belgium (2016), when terrorists, trained in Syria, murdered more than 200 civilians in areas otherwise thought secure and gained the Islamic State of Iraq and the Levant (ISIL) a strategic information victory (Almasy 2015, Shoichet 2016).

Figure 2. Foreign Fighter Flow into Syria. Adapted from Sharma (2015).

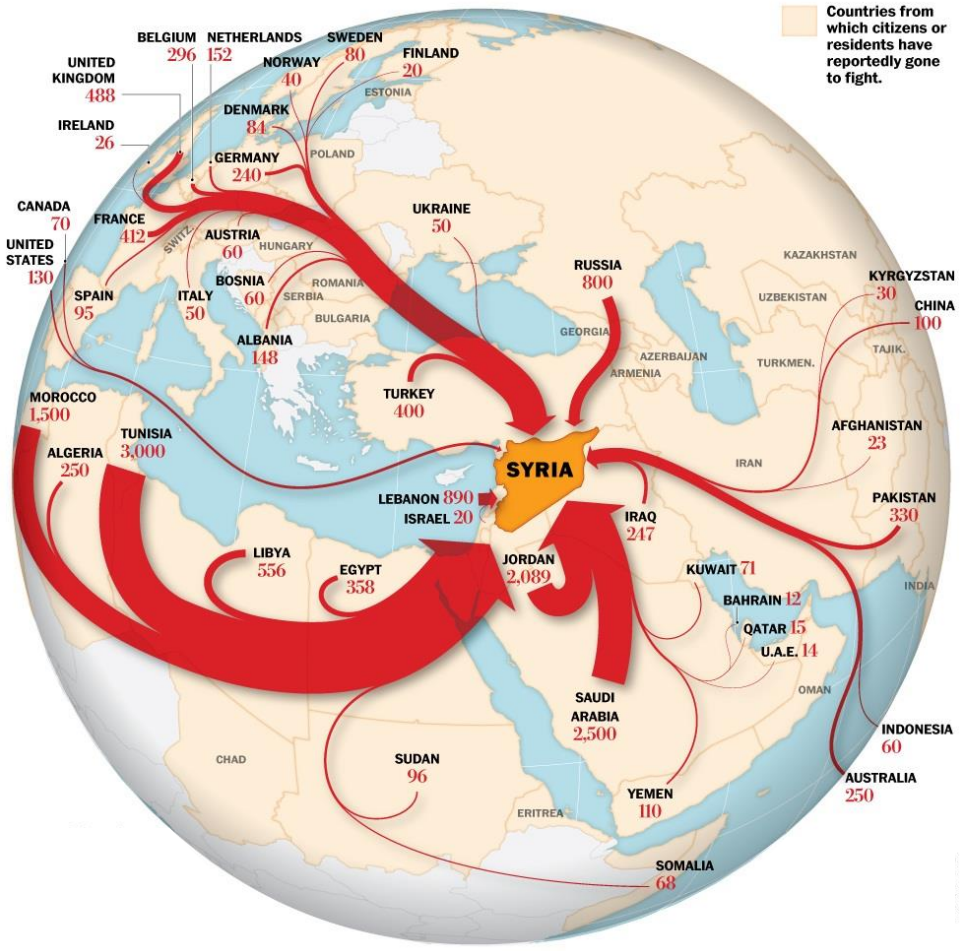

While it is unlikely that the flow of fighters, materiel, or terrorists can be completely stopped, more well-coordinated interdiction efforts can significantly influence smuggling behaviors and potentially seize key shipments. These behavioral changes and the threat of interdiction could help drive criminal smuggling organizations from cooperating with terror groups, further limiting the scope and impact of either's activities.

The study of this interdictor-smuggler contest is important, yet difficult. There is no single, natural methodology to address this scenario with the features we have suggested, and it does not gracefully decompose into sub-processes that submit to isolated analysis and produce meaningful results. As a result, making robust statements of cause and effect is problematic in these circumstances. However, by using a family of analytical techniques, it is possible to comment on emergent patterns that reflect the situation's time dynamics and eventual rest points.

B. PROBLEM STATEMENT AND SCOPE

In this thesis we consider a resource allocation problem faced by the interdictor, specifically, where to position limited combat power along suspected infiltration routes over time. This must be done using incomplete information against an adaptive smuggler. The desired objective is to best-disrupt the smuggler's lines of communication by minimizing the amount of materiel or goods travelling from source to target.

Our approach is bottom up. We merge strong aspects of traditional game theory, optimization, stochastics, and simulation. First, we design a game with simple rules to examine patterns of interest and search for recognizable emergent behaviors. This game is instantiated in a terminating discrete event simulation. By repeating this simulation in a variety of configurations, each continued over a finite number of time steps, we study the time dynamics of our problem in a spectrum of situations. Through this evolution of play, we examine the robustness of interdictor resource allocation policies and smuggling tactics. Lastly, we search for emergent relationships during play.

Our computational study begins with traditional relaxations that include perfect information for the interdictor and smuggler. We then demonstrate the effect of private information for each player. Both the interdictor and smuggler must make estimations and decisions because of the partial information on their antagonist's current state and designs. Private information and estimation introduces adaptive play under uncertainty. We assess the value of this feature. Next, we simulate games configured with several specific interdictor and smuggler policies. Finally, we examine each game's time dynamics and end-results through the lens described above.

We have several goals. First, we aim to expose practical insights for decision makers within a counterinsurgency or counter-illicit trafficking setting. Our study should highlight otherwise-obscure relationships between interdiction policy and outcome while identifying realistic ways to increase the effectiveness by which resources are allocated. We expect the results could bring attention to scenarios that military decision makers might neither be able nor want to explore, due to a lack of information and human bias, respectively. Our formulation and results might suggest the importance of many features otherwise absent from current models. Some implications of our model may conflict with prior beliefs, provoking new questions and lines of research inquiry. Lastly, we strive to demonstrate an analytic way to investigate an important, realistic, and yet intensely complex problem by rigorously blending several complimentary stochastic, game-theoretic, and optimization models.

II. LITERATURE REVIEW

The key features underlying the interdictor-smuggler problem have been given significant treatment in both United States military doctrine and Operations Research. First, we explain the nature of war as defined in modern U.S. warfare philosophy. Then, we examine the attributes describing the essence of conflict: uncertainty, fluidity, and adaptation. Next, we review the analytical models that address interdictor-smuggler problems. In the survey of analytical techniques, we relate each model's assumptions to the attributes of war as described in U.S. military doctrine. Lastly, we outline the contribution of this thesis.

A. MODERN UNITED STATES WARFARE PHILOSOPHY

Modern United States military doctrine begins by defining the essence of war. The agreed upon definition is: "a violent struggle between two hostile, independent, and irreconcilable wills, each trying to impose itself on the other" (Marine Corps Combat Development Command 1997). A common view of conflict helps aid interoperability and sets the foundation for all further service-level discussions of operations and tactics. *The Joint Operating Environment* leverages the agreed essence of war and "provides a perspective on future trends, shocks, contexts, and implications" for the near and far term security environment facing the United States (U.S. Joint Forces Command 2010). Its key points expound on the essence of war and how its precipitates—uncertainty, fluidity, and adaptation—will continue to govern the course of events in conflict. Each military service devotes extensive discussion throughout their respective capstone doctrine and tactical publications to defining and addressing the importance of uncertainty, fluidity, and adaptation (e.g., United States Air Force 2003, Department of the Army 2012, Marine Corps Combat Development Command 1997, U.S. Navy Doctrine Command 1995).

Uncertainty is the prime attribute of war in all U.S. Military doctrine. The U.S. Air Force explains that the incompleteness of information is so pervasive on the battlefield that it permeates combatants' views of the enemy, environment, and even their

own forces (United States Air Force 2003). Elder (2006), in a study by the Central Intelligence Agency, concludes that the accuracy of assimilated information was a decisive factor in five strategically significant battles: the First Battle of Bull Run (1861), Tannenberg (1914), Midway (1942), Inchon (1950), and the Israeli air strike initiating the Six-Day War in 1967. Each Service argues that uncertainty can never be eliminated (e.g., Department of the Army 2012). Because of this intractability, the "fog of war" relates directly to another concept, fluidity.

Fluidity communicates the tension of two competing phenomenon, uniqueness and dependence. U.S. Marine Corps doctrine establishes that while each combat engagement is a unique composition of circumstances, it is also dependent on the myriad events before it and determines those engagements that follow it (Marine Corps Combat Development Command 1997). The U.S. Army (2012) agrees with the importance of examining the path-dependency of each combat, citing that no combat episode can be viewed in isolation. They also sustain that no combat episode repeats itself exactly (Department of the Army 2012). Each engagement thus requires an original solution according to the U.S. Naval Doctrine Command (1995). U.S. Military doctrine concludes that because of uncertainty and fluidity, combatants must adapt.

Adaptation is recognized by U.S. military doctrine as an imperative to success in war. Colonel John Boyd, the inventor of the *Observe-Orient-Decide-Act* decision framework, describes adaptation in combat through the lens of *orientation* (Boyd 1976). A combatant's current awareness and experiences define their orientation, full of uncertainty and prejudice. Events not anticipated by this orientation generate surprise. Colonel Boyd explains that success in conflict is based primarily on each belligerent's ability to anticipate or recognize these anomalies and then reconcile them quickly and accurately (Boyd 1976). The proliferation of Colonel Boyd's theories during the 1980s brought adaptation back into the center of modern military doctrine. His influence is readily apparent in Marine Corps Doctrinal Publication 1 (1997), the Marine Corps' capstone doctrinal publication. It defines adaptation and reinforces the context of uncertainty in conflict:

> War is thus a process of continuous mutual adaptation, of give and take, move and countermove…The very nature of war makes certainty impossible; all actions in war will be based on incomplete, inaccurate, or even contradictory information. (3-4)

The doctrine and tactics manuals from the other Services also give extensive treatment to the importance of adaptation (e.g., U.S. Department of the Army 2008). Even today, adaptation remains a centerpiece of martial discussions.

The authors of U.S. military doctrine, in reviewing the recorded history of human conflict and leveraging decades of combat experience, communicate the supreme importance uncertainty, fluidity, adaptation, and complexity play in determining the outcome of the activities surrounding war. This closely aligns with the author's own combat experience as a counterinsurgent.

B. OPERATIONS RESEARCH

The field of Operations Research arose in World War II to provide quantitative decision support to both U.S. and UK commanders within the wartime context of uncertainty and complexity described above. Its analytical methods were successfully applied to a spectrum of problems including convoy protection, amphibious assault, and aerial bombardment (Kirby 2003). Project Research ANd Development (RAND) was created during this time by General Henry H. Arnold, then commander of the United States Army Air Forces (RAND 2016). It was one of the major efforts by the U.S. to leverage quantitative analysis for decision support in World War II. The genesis of modern interdiction studies occurred within Project RAND.

1. Cold War Origins

On 21 May 1999, the United States Air Force declassified a study originally commissioned under Project RAND in 1955 (Harris and Ross 1955). It revealed the first modern mathematical model that we would recognize as network interdiction. The study describes a mathematical problem to identify a set of most vital routes within a transportation system whose removal would completely stop all movement. This problem has been extended in a number of directions, now commonly describing a contest

between two completely opposed and intelligent adversaries seeking, respectively, to inhibit or enhance the flow from one or more source nodes to one or more target nodes within a network. An explosion of additional interest since the 1990s has resulted in myriad techniques that now attempt to treat both a variety of contexts and assumptions through several related general formulations. (For further information, see Alderson et al. [2013] for a detailed treatment of this problem and references to subsequent work.)

Harris and Ross advocated a basic model now referred to as *k-most vital arcs*. In the context of a potential Soviet invasion of Western Europe, they sought to stop the flow of reinforcements within the Soviet railway network. These forces and materiel would transit from sources within the Eastern Soviet Union to destinations in the western reaches of the Soviet Union and her satellite states in Eastern Europe. Recognizing that contemporary methods were inadequate, they instead postulated a holistic view that considered both primary and alternate routes and aggregated railway-operating divisions. Their solution identifies a "bottleneck," a set of arcs that, if cut by aerial strikes, would completely interdict all western flow (Figure 3).

Figure 3. The Contemporary Railway System of the U.S.S.R. with Identified "Bottleneck" as depicted by Harris and Ross. Source: Harris and Ross (1955).

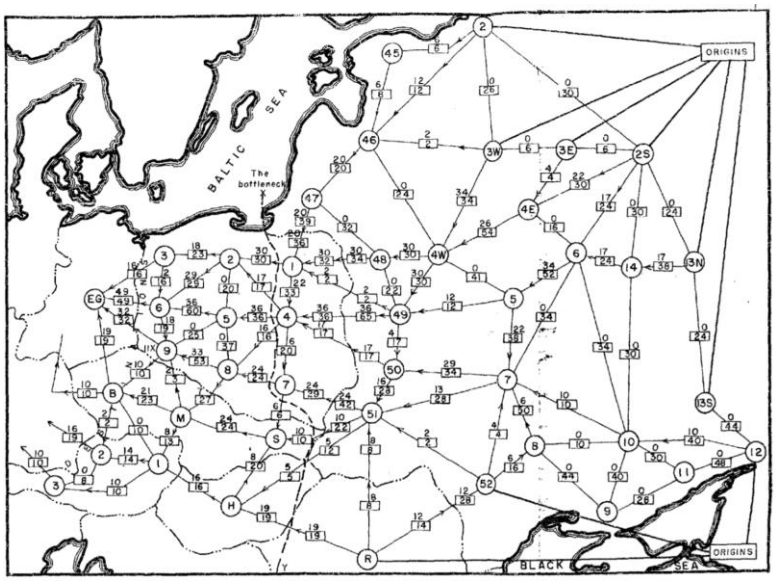

Nodes represent aggregate railway operating divisions and arc information represents the capacity in thousands of tons that can be moved between divisions in one day.

Importantly, Harris and Ross (1955) comment on two aspects of the problem that must be considered:

> As in many military operations, however, the success of interdiction depends largely on how complete, accurate, and timely is the commander's information, particularly concerning the effect of his interdiction-program efforts on the enemy's capability to move men and supplies. (iii)

> …it is fully recognized that the difficulties inherent in obtaining, evaluating, and disseminating intelligence would limit the usefulness of the method in direct proportion to the information placed at the disposal of the particular specialist concerned. (2)

We argue that in many martial circumstances, such as the interdictor-smuggler scenario we posit, Harris and Ross have identified an intractable factor of these problems, not an inconvenience that can be overcome in assumptions. A number of interdiction models have addressed this issue directly or made strong assumptions concerning it for tractability.

2. Basic Models

There are several basic types of interdiction models. First, Harris and Ross (1955) propose to estimate the maximum railway network flow by identifying and then calculating the capacity of limiting bottlenecks, as described above. In 1963, Wollmer followed the work of Harris and Ross, describing a model that maximally reduces flow through a rail system by finding and cutting the *most vital arc*. Corely and Sha (1982) extended this model to consider weighted arcs and nodes but proposed only binary interdiction of identified arcs to minimize the maximum possible flow. In this sense, arcs are either completely cut or left uninhibited. Fulkerson and Harding (1977) adopt another model, *shortest path interdiction*, seeking to maximize the minimum *source-target* (*s-t*) path subject to a budget constraint. These basic ideas have been merged by Malik et al. (1989), where the k-most vital arcs are interdicted with the goal of maximally increasing the length of the shortest *s-t* path. Israeli and Wood (2002) formulate this method as a bi-level program, introducing "supervalid inequalities" and greatly decreasing computational time. The third basic model is called *maximum flow interdiction*. Under

this perspective, arc costs or capacities are degraded to minimize the maximum possible flow (Wollmer 1964). In the shadow of the Vietnam War, these models were applied most directly to attacking enemy logistics systems (Ghare 1971, Ratliff et al. 1975). Importantly, these basic models assume that both interdictor and evader have perfect knowledge of both the environment and one another's range of potential actions.

Competing, intelligent adversaries have been an implicit feature of all of these models. Typically cast as two-person zero-sum games, network interdiction problems closely align with game theory. Danskin (1966) establishes the foundation of a generalized theory and solution method for these max-min problems. For exposition, he proposes a situation in which a defender installs fortifications and then an attacker plans a strategy with full knowledge of the location of these installations. This kind of sequential gaming format is closely related to a Stackelberg game and commonly referred to as such within modern network interdiction literature (von Stackelberg 1952). See Wood (2011) and references therein for a thorough discussion. Simultaneous gaming constructs have also been applied to network interdiction models, but this is less common. Washburn and Wood (1995) addressed the interdiction of illicit drugs by method of maximum flow interdiction and simultaneous gaming. While these models describe a complex problem, they still describe one in which the possible state space is known with certainty.

3. Introduction of Uncertainty

Natural extensions of these models include the introduction of uncertainty. Stochastic network interdiction models were first introduced by Cormican et al. (1998). Here the success of an interdictor's attacks is binary but uncertain. Extensions are described that include uncertainty in arc capacities but require certainty in attack outcome (Cormican et al. 1998). Several efforts in support of counter-nuclear smuggling have formulated and solved stochastic network interdiction problems (Morton et al. 2007). In these problems, arc costs represent detection probabilities. The interdictor's goal is thus to maximize this probability across all possible infiltration paths given an unknown smuggler origin. Using a logarithmic transformation, the model then becomes deterministic. Even so, the model is stochastic in the whole because the smuggler's origin

is described by a probability distribution. Nehme (2009) extends this model by examining sequential, simultaneous, and then hybrid games where only a portion of the radiation sensors are made visible to the smuggler. Further relaxations are often noted as intractable under the above formulations.

These stochastic programming models indirectly include asymmetric information. Several recent studies have explicitly made this inclusion but maintained static network character and do not consider time dynamics. Uniquely, Salmeron (2012) explores deception tactics under network interdiction. He formulates a multi-objective bi-level program that optimally locates a number of covert sensors, overt sensors, and decoys. As in the counter-nuclear smuggling models above, the focus of this interdictor is to maximize the probability of detection for a single evader. He notes that the computational time to reach solutions within this model is extensive.

Asymmetry of information has been treated extensively in game theory literature. The literature centers on a problem of exploration and exploitation, commonly now referred to as the *multi-armed bandit problem* (Robbins 1952). Within this problem, a gambler is faced with a number of slot machines, each of which produces outcomes drawn from a unique distribution. The gambler must then decide how to balance *exploration* by playing various machines to forecast which will provide the highest reward and actually *exploiting* this information by playing this subset of machines. The aim is to maximize the sum of rewards. More nefarious versions have been proposed that pit the player against a malicious casino that controls the game payoffs (Auer 1995).

These multi-armed bandit problems attempt to confront directly the military intelligence problem described by Harris and Ross and implicitly linked to practical network interdiction. Coping with initially incomplete information that can be uncovered in time requires feedback mechanisms. This presents problems with dynamic instead of static data.

4. Dynamic Data

Zinkevich (2003) introduces online optimization for repeated games by convex programming in an attempt to address the challenge of dynamic data. This method can be described by a procedure that is performed in every time step:

1. Choose an action.
2. Simultaneously an adversary selects an action.
3. Suffer loss that is a function of both selected actions.
4. Observe the adversary's action.

The objective is to minimize the cumulative loss over time (Bubeck 2011). In this protocol, feedback is perfect, or *transparent*; the player is informed of the adversary's complete actions without error. Awerbuch (2004) addresses a minimum delay routing shortest path optimization problem with *opaque* feedback. Here a malicious and adaptive adversary's actions only partially reveal the underlying network structure. True arc costs are made visible just on selected paths. Similar formulations also provide the forecaster with a subset of the true loss vector (Cesa-Bianchi et al. 2012). It is apparent that these models attempt to confront more directly exploration-exploitation issues that emerge in dynamic optimization problems.

Online stochastic optimization convolutes the problem by further limiting feedback. One of the goals of the limited feedback models is to explore the strategies required to cope with partial information scenarios. Bubeck (2011) defines *bandit feedback* as feedback wherein the player observes the adversary's moves only indirectly, confounded with other factors. The *semi-bandit* version allows perfect loss information, but only in areas explored by that turn's active strategy (Bubeck 2011).

A most recent effort executes network interdiction across time with feedback. Borrero et al. (2015) address sequential shortest path interdiction with partial information. The interdictor has incomplete knowledge, but the evader has complete knowledge. As the evader traverses the network turn-by-turn, additional arcs and accurate costs are revealed to the interdictor through semi-bandit feedback. The authors suggest assessing interdictor policies by *time stability* and efficiency. This time stability is the

number of time steps required before chosen strategies match those of a player with perfect knowledge called the *oracle* (Borrero et al. 2015). In this case, the adversary is myopic and non-adaptive.

With the rise in computational power and development of newer decomposition methods, modern network interdiction studies are now able to handle what would have been prohibitively large problem instances in the time of Harris and Ross. However, reflecting again on their admonitions concerning imperfect information and time sensitivity, we can see that modern techniques still require a multitude of strong assumptions that ignore the important effects of two-sided asymmetric information, path-dependency, and adaptation. We argue that these assumptions have played a role in limiting both the scope of current network interdiction research and practical utility of some results.

C. OUR CONTRIBUTION IN CONTEXT

We propose a novel network interdiction model that reconciles many operational realities identified by military literature. We believe that significant insight can be gained into heretofore underrepresented or excluded problems within these operational realities. We pursue this insight by simultaneously relaxing many of the previous network interdiction modelling assumptions listed above. The local counterinsurgent-smuggler contest described in Chapter 1 is one example in this class of problems. For this problem, we conduct network interdiction within a dynamic network, under partial information, using incomplete feedback, and allowing two-sided adaptive play. We combine these aspects in an evolutionary game, leveraging optimization, simulation, and stochastics to achieve a hybrid model.

THIS PAGE INTENTIONALLY LEFT BLANK

III. MODEL FORMULATION

In this chapter, we describe a game and corresponding mathematical model formulation for gaining insights into realistic smuggling and interdiction scenarios. First, we define the objectives, rules, interactions, and challenges that form the game. Next, we explore a series of constructive modelling cases to demonstrate the effect of each major modelling assumption. We then explain the mathematical formulation of our game.

A. THE GAME

We design a game in which two players with opposing goals make decisions that govern the movement of materiel across a network. The game is played in discrete periods called *rounds.* The goal of the first player, named the *smuggler*, is to move as much of this materiel as possible from one or more *sources* to a single *target*. Opposing the smuggler is the second player, called the *interdictor*, who attempts to stop this materiel from reaching the appointed target. Each round of the game produces a *score* that represents the total materiel delivered to the target in that round of play. Playing the game under different scenarios allows for the relative evaluation of various tactics and provides insights into phenomena modelled by the game.

The game proceeds over a finite number of rounds. Neither player knows the total number of rounds within a game. Each round involves several decisions by both the smuggler and interdictor. While these two players make some decisions in turn, they also make some decisions simultaneously. The outcome of a player's decisions is stochastic. Each player updates the information used to make their decisions based on these results. Score is recorded. The game then proceeds to the next round.

The smuggler and interdictor play this game on a network described by a directed graph consisting of *nodes* and *arcs*. Nodes represent an origin, destination, or intermediate location for materiel flow. Thus, both the source and target are nodes. There can be multiple sources within a game and even within a single round of play but only one target. Nodes have unbounded storage capacity. Arcs represent potential movement of flow from one node to another node. There is only one arc between any pair of nodes,

and the underlying graph is acyclic. These two characteristics make the network a *feed forward* network. Each arc has a unique *cost* and *capacity*. The cost represents the amount of work required to move one unit of materiel from the node at the tail of the arc to the node at the head of the arc within a single round of play. Similarly, the capacity describes the maximum amount of materiel that can be moved within a round of play across the respective arc.

Materiel moves in discrete units called *packets*. Each packet has two attributes: a source and a *size*. The size of the packet describes the amount of materiel contained within the packet. Different packets can have different sizes, but the size of each packet is fixed. A packet's source and the target determine the node in which the packet enters the game and the node from which the packet can exit the game, respectively. Packets reaching the assigned target node within a round count toward the score. The round's score is tallied by summing the sizes of these packets.

Both the smuggler and interdictor play the game by making different decisions in the context of individual budgets. The smuggler decides how to use a finite budget to move the packets through the network. This budget, the *movement budget*, can change by game round. Each packet moves as a unit. Within a round, the length of movement may be limited to one arc or extend to the entire set of arcs connecting the source and target nodes. Alternatively, the smuggler may decide not to move some packets at all within a game round.

The interdictor decides where to place a limited number of sensors to discover the packets and prevent them from reaching the target. A *sensor budget* limits the number of sensors the smuggler may place in a game round. As with the smuggler, this budget can change by round. Once placed, we assume a sensor lasts only one round. Sensors are either *overt* or *covert*. Overt sensors are visible to the smuggler while covert sensors are not. Both types of sensors are placed upon arcs. The interdictor may place only one sensor of any type upon any one arc in any one round of the game. Packets passing over an arc on which a sensor is placed are candidates for detection. This detection is random and the probability of detection is a function of the packet's size and the number of past detections on the arc. Even so, the interdictor does not know the probability of detection.

The interdictor automatically attacks detected packets, destroying them. We remove destroyed packets from the game.

The information of both the smuggler and interdictor is always incomplete. Even so, the information is incomplete in fundamentally different ways. The smuggler lacks complete awareness of all available arcs, and the interdictor is uncertain of all arc and node attributes. Both players formulate estimates to augment this partial information. The smuggler's estimates of arc costs and capacities are equal to the ground truth arc costs and capacities. However, the set of arcs visible to the smuggler is a subset of the ground truth arcs. The visibility of an arc is controlled by a clock that reveals the arc to the smuggler in a predefined round. After an arc becomes visible, it is always visible.

The interdictor also has an estimate of network information and smuggler decisions. A limited feedback loop informs the interdictor's sensor placements by formulating a private view of arc capacity, arc cost, and the location of the packets. Because the estimated locations and target of the packets—the estimated node *supplies* and *demands*—shifts round-to-round, and the interdictor is attempting to prevent the smuggler's flow, the interdictor considers different arcs in each round. Through feedback loops, "learning," and "forgetting" occurs. Each player then adapts to better inform their play and commit resources.

We apply a *transformation function* to all the arcs within the game at the end of each round. This function reduces each arc's cost by a fractional amount every round for three purposes. First, the transformation function implements the count-down timer to reveal non-visible arcs to the smuggler. Second, the transformation function simulates smuggler learning by decreasing the cost to transport materiel in successive rounds. Third, the transformation function attenuates the smuggler's cost increase resultant from any loss of materiel incurred during previous game rounds. The combination of the transformation function and evolving player's estimates gives the network a dynamic character.

We consider numerous *scenarios* under which the interdictor and smuggler play the game, in order to evaluate a spectrum of smuggler and interdictor policies. Modifying the sensor budget, movement budget, or other parameters can cause each player to use

different tactics and allow one to study the effectiveness of a number of styles of play. It is then possible to examine the relative value of each tactic and make observations on the progression of play within the game in order to draw more general conclusions on phenomena that might be well-represented by this game.

B. THE I-I/S SIMULATION ALGORITHM

We develop an heuristic algorithm to simulate the time-ordered decisions made by the interdictor and smuggler in the game described in Section A (Figure 4). Within the algorithm, we formulate the game as two-stage sequential: *Interdictor – Interdictor/Smuggler* (**I-I/S**). In the second step of the round, "*Interdictor/Smuggler*," play is simultaneous. The interdictor and smuggler play the game over a finite number of rounds.

The algorithm allows us to explore the performance of both the interdictor and smuggler under various individual resource allocation policies. We instantiate the algorithm in a terminating discrete event simulation. Evaluation of the output provides feasible, face valid solutions to various interdictor-smuggler problem instances.

Figure 4. The **I-I/S** Simulation Algorithm.

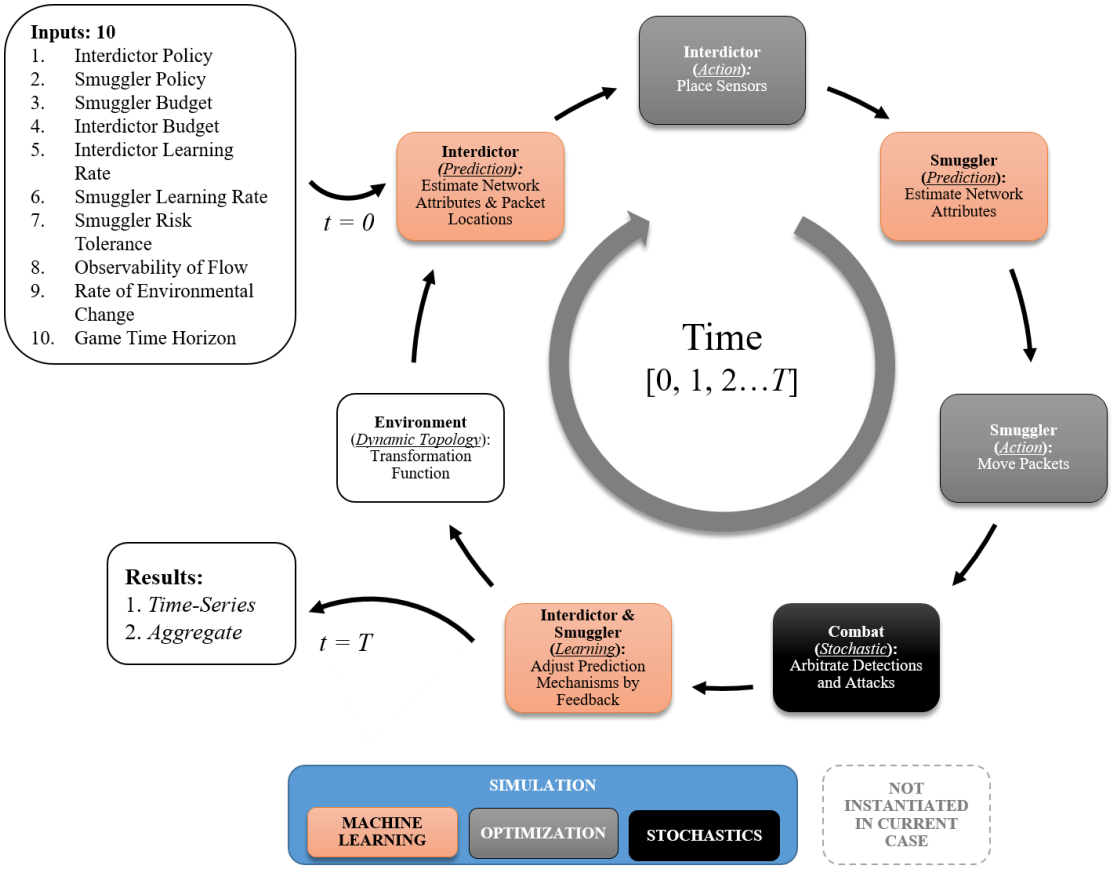

The time-ordered steps of the **I-I/S** game algorithm. Nine initial inputs set the specific scenario. At each step in the algorithm, one or both players makes decisions based upon or adjusts a prescriptive model that represents their current state of knowledge about the system and the actions (so far) of their opponent, and explicitly models any limitations of this information and any uncertainty in the outcome of their actions. After a pre-determined number of game rounds *T*, the algorithm generates both time-series and finite time horizon aggregate outputs.

C. CONSTRUCTIVE CASES

We explore a series of constructive cases to demonstrate the effect of each major modelling constraint (Table 1). Each constructive case considers the same problem instance involving six nodes, two of which are sources and one of which is the target for the smuggler (Table 2 and Figure 5). We use these cases as illustrative examples to further explain game play and to argue that the admixture of the constraints found in our full model is necessary to properly explore the interdictor-smuggler resource allocation problem.

Table 1. List of Constructive Cases.

	Case 1	Case 2	Case 3	Case 4	Case 5
Multi-period	X	X	X	X	X
Asymmetric cost and capacity information		X		X	X
Limited cost and capacity feedback		X		X	X
Packet locations unknown to interdictor			X	X	X
New arcs revealed by timer					X

Table 2. Parameters and Values for the Problem Instance Simulated in Cases 1–5.

Interdictor Budget (sensor budget)	1 Overt Sensor and 1 Covert Sensor / round
Smuggler Budget (movement budget)	25 / round
Packets	3 packets / round
Number of round played	3 rounds / Case

Figure 5. Network Configuration and Packet Schedule for Constructive Cases.

Network Topology

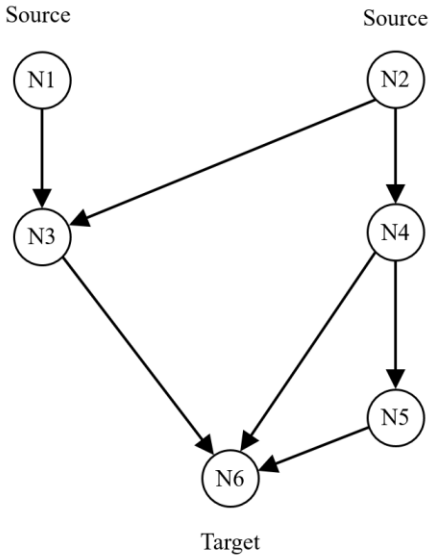

Packet Schedule

Round	Packet	Location
1	P1	N2
1	P2	N2
2	P3	N2
2	P4	N2
3	P5	N2
3	P6	N1
3	P7	N1

Left: The network configuration for all constructive cases. Nodes *N1* and *N2* act as sources, node *N6* acts as the target. *Right*: The schedule of packets for all computational cases. The schedule introduces seven total packets, all of *size* 1, throughout the game. Note that in round three, the schedule introduces two packets at node *N1*.

1. Case 1: Multi-period, Symmetric Information

In Case 1, we assume both the smuggler and interdictor observe the same arc costs and capacities. The interdictor has perfect information of the system state. All packet locations are transparent. The smuggler can see overt sensors but is not able to see covert sensors. Because the smuggler can neither observe nor predict the placement of covert sensors, he moves packets in a greedy, myopic fashion along budget feasible paths. The entire set of arcs is visible to the smuggler throughout the game. Figure 6 displays the time-ordered steps of the game algorithm for constructive case 1. Figures 7–9 display the decisions, information, and outcomes of three rounds played under Case 1.

Figure 6. The **I-I/S** Game Algorithm, Case 1.

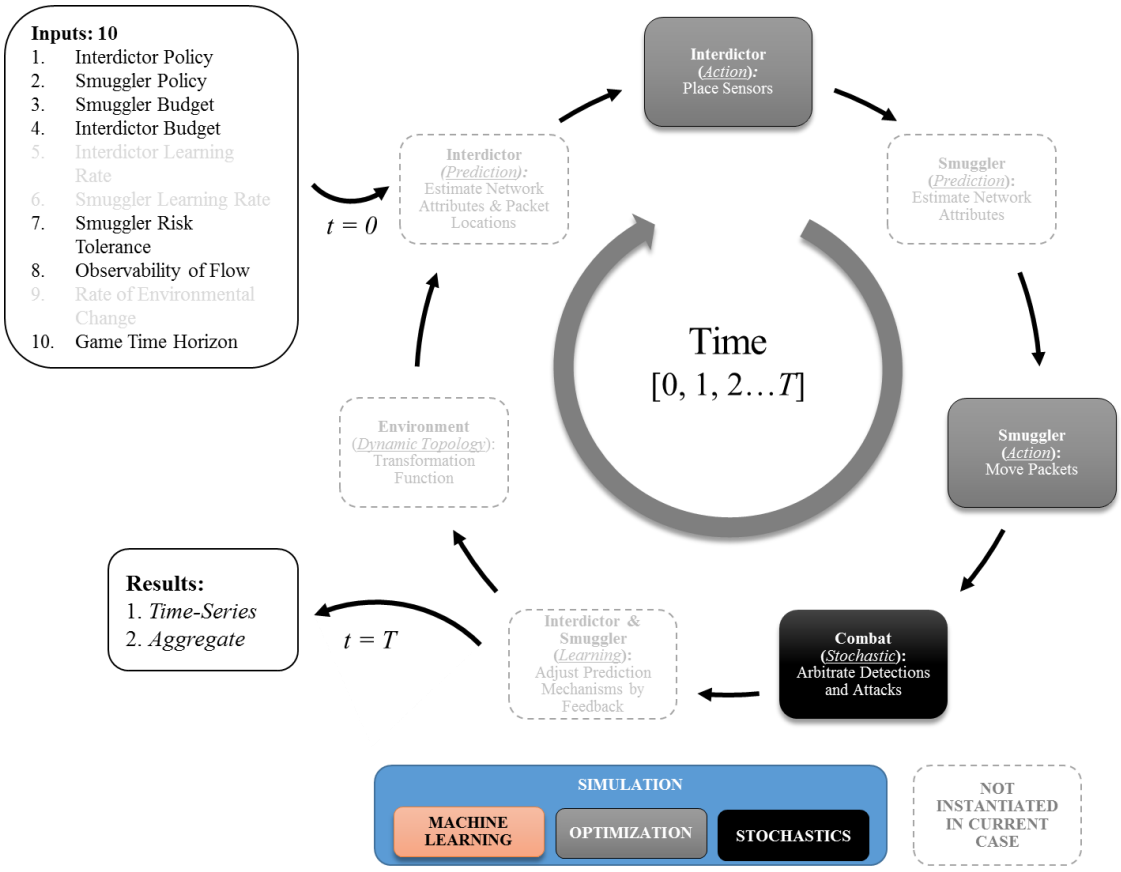

The time-ordered steps of the game algorithm for constructive Case 1. In Case 1, the interdictor has perfect information of the system state, requiring neither estimation nor learning on his part. The smuggler is able to view all arcs within the network during all rounds of play. Detection success is stochastic. The Case 1 configuration requires only three of the **I-I/S** game algorithm's seven steps to instantiate (unnecessary steps in dashed gray).

Figure 7. Case 1, Round 1 Decisions and Information.

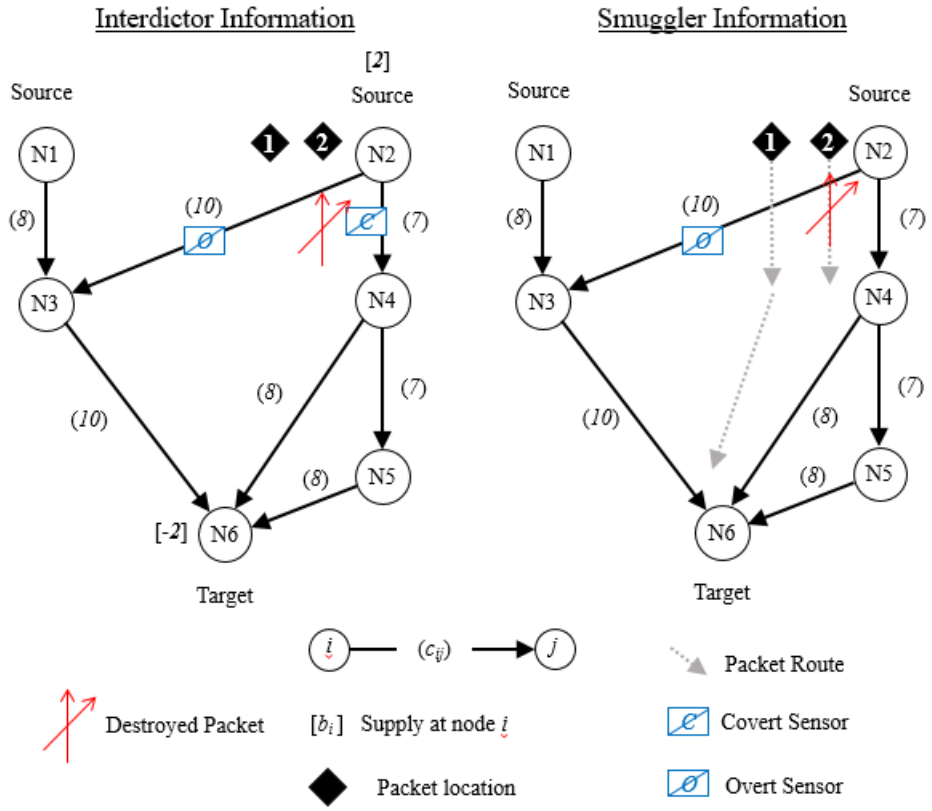

The smuggler has two packets at node *N2*, and the interdictor is aware of all packet locations. *Left (Interdictor)*: The interdictor optimally places two sensors to maximize the smuggler's minimum cost of flow. *Right (Smuggler)*: The smuggler is aware of the interdictor's overt sensor on arc (*N2*, *N3*), but unaware of the covert sensor on arc (*N2*, *N4*). He moves two packets to the target at minimum cost based on his partial information (dashed). The stochastic result from the two packets passing the covert sensor on arc (*N2*, *N4*) yields "no detection" for packet 1 but "detection" for packet 2. The interdictor destroys the detected packet. Packet 1 reaches the target and records a score of 1.

Figure 8. Case 1, Round 2 Decisions and Information.

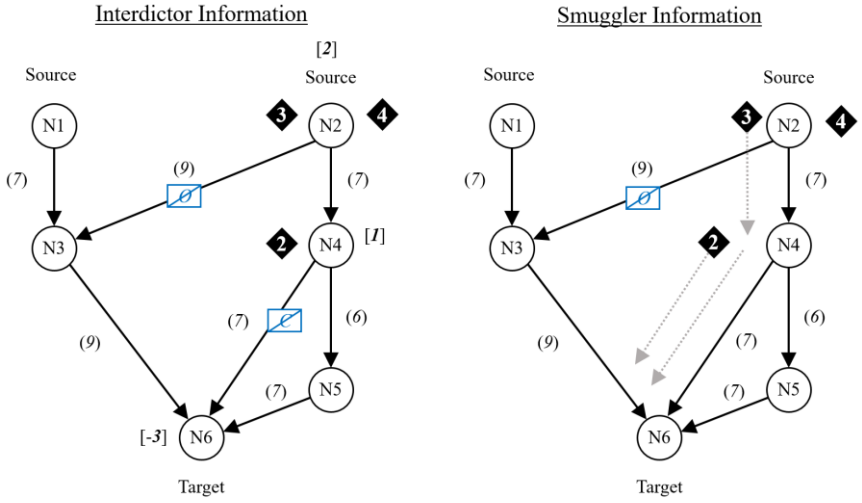

Packet 2 is at node *N4* from the previous round, and there are two new packets at source node *N2*. The transformation function updates all arc costs, reducing them. *Left (Interdictor)*: Using the same arc attributes as the smuggler in Case 1, the interdictor again places sensors optimally. *Right (Smuggler)*: The smuggler attempts to move packets 2 and 3 to the target. The stochastic result from these packets passing the covert sensor on arc (*N4, N6*) yields "no detection," and a score of 2 is recorded.

Figure 9. Case 1, Round 3 Decisions and Information.

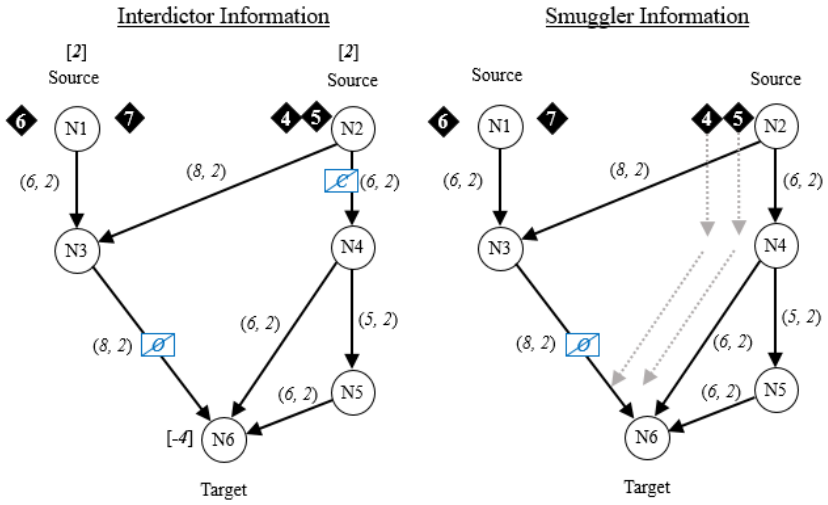

Packet 4 remains from the previous round, and there are three new packets at source nodes. The transformation function reduces all arc costs. *Left (Interdictor)*: Aware of all supplies and arc costs, the interdictor optimally places sensors to maximize the smugglers minimum cost. *Right (Smuggler)*: The smuggler attempts to move two packets to the target. Because of the overt sensor on arc (*N3, N6*), the smuggler does not send packets 6 and 7, and instead sends packets 4 and 5. The stochastic result from packet 4 and 5 passing the covert sensor on arc (*N2, N4*) yields "no detection." Both packets reach the target giving a score of 2.

Discussion

In Case 1, we assume the interdictor is aware of all packet locations. He, therefore, places sensors as closely as possible to these known supplies of flow in each round. Additionally, because the smuggler and interdictor use the same cost and capacity information, the interdictor has a perfect prediction of the smuggler's packet routing. Sensor detection probability is less than 1.0, so even with the completeness of information available, five packets still reach the target. The model in Case 1 is very similar to previous stochastic network interdiction models where perfect information exists, but attack success is uncertain (e.g., Cormican et al. 1998). The assumption of perfect information made in Case 1 does not comport well with the key features of the interdictor-smuggler problem as posed in Chapters 1 and 2.

2. Case 2: Multi-period, Asymmetric Cost and Capacity Information, Limited Feedback

In Case 2, we assume the smuggler and interdictor have private arc cost and capacity information. Additionally, we suppose that the private cost information is substantially different for illustrative purposes. All packet locations are still transparent. The interdictor must make a prediction on the smuggler's movements by forming estimates of the arc costs and capacities. Indirect, limited feedback updates the information of both the smuggler and interdictor. The smuggler is aware of overt sensors, but is never made directly aware of the location of covert sensors. Instead, the smuggler is aware of packets destroyed by the interdictor. The smuggler uses these losses to estimate heightened detection risk (or threat locations) by increasing his estimate of arc cost on the arcs where the interdictor destroyed packets. In similar fashion, successful detections reduce the interdictor's estimate of arc cost and tune his estimate of arc capacity. Arc capacity must be at least equal to the amount of materiel detected in a single round on the arc. However, the interdictor knows the smuggler's supply nodes because the packet locations are transparent in this case. Lastly, the entire set of arcs is visible to the smuggler throughout the game. Figure 10 displays the time-ordered steps of the game algorithm for constructive Case 2. Figures 11–13 display the decisions, information, and outcomes of three rounds played under Case 2.

Figure 10. The **I-I/S** Game Algorithm, Case 2.

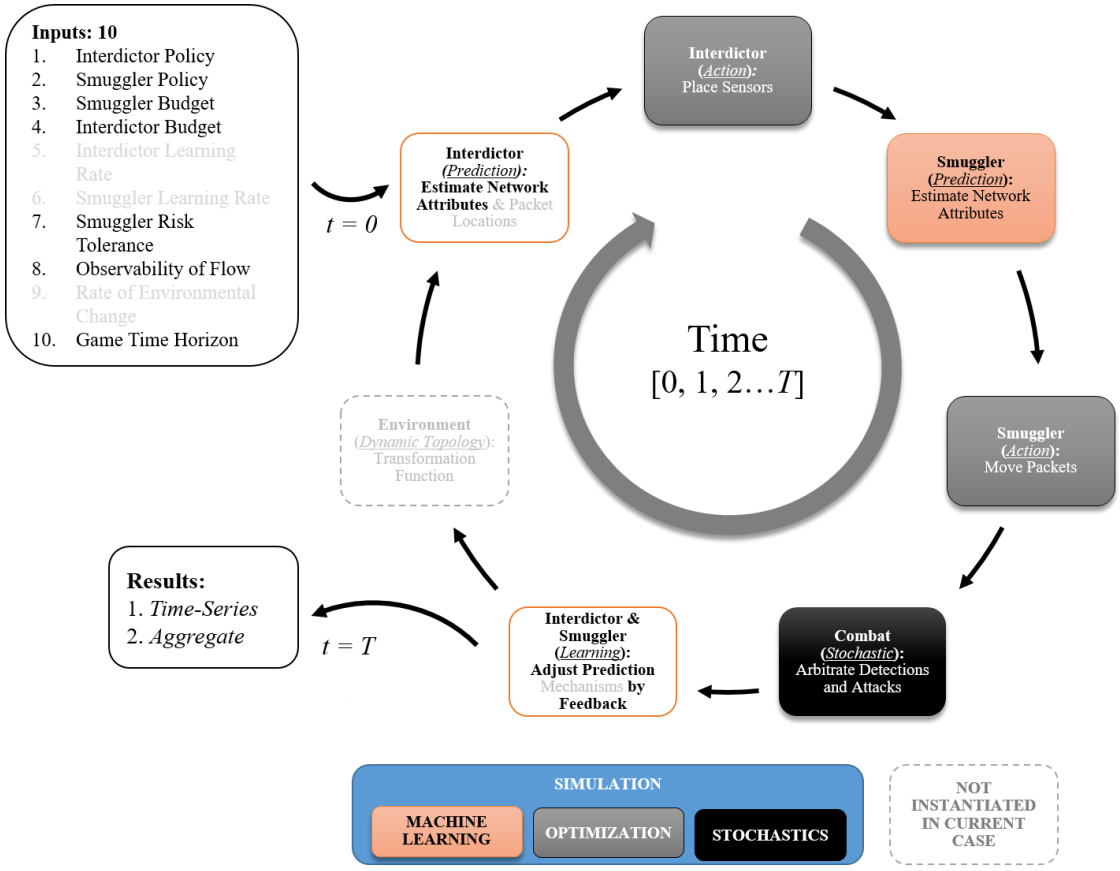

The time-ordered steps of the game algorithm for constructive Case 2. In Case 2, we introduce asymmetric information. The information asymmetry requires the interdictor to estimate the arc attributes. Even so, the interdictor is aware of all packet locations. Only part of the "Interdictor (Prediction)" step is thus required. Both the smuggler and interdictor update their incomplete information by limited feedback. However, only one of two feedback mechanisms is in place for the interdictor.

Figure 11. Case 2, Round 1 Decisions and Information.

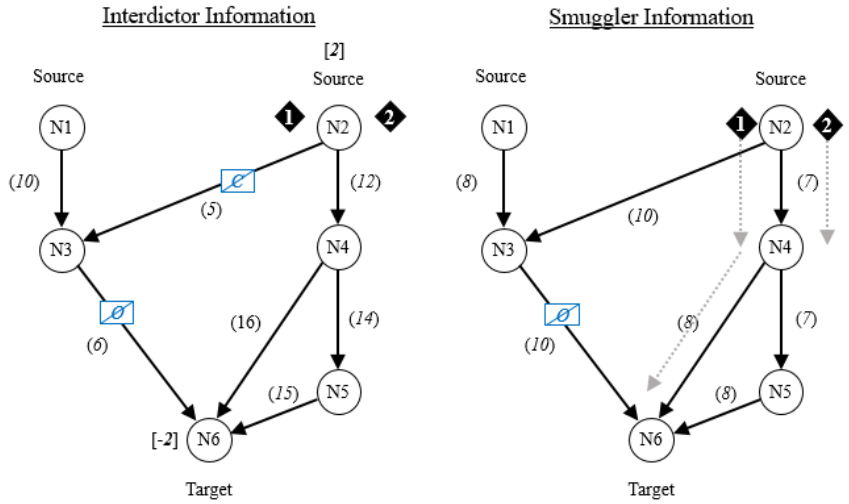

This case follows the same packet schedule as before, starting with two packets at source node *N2*. *Left (Interdictor)*: In Case 2, the interdictor must make his own estimate of arc attributes. For illustrative purpose, we suppose that these estimates are initially substantially different from the smuggler's information. Even with perfect information on the location of all packets in play, the inaccuracy of the interdictor's initial estimate causes the him to sub-optimally place a covert sensor on arc (*N3, N6*) instead of arc (*N2, N4*). *Right (Smuggler)*: As in Case 1, the smuggler is unable to see covert sensors. He moves two packets optimally given his partial information. Packet 1 reaches the target and records a score of 1.

Figure 12. Case 2, Round 2 Decisions and Information.

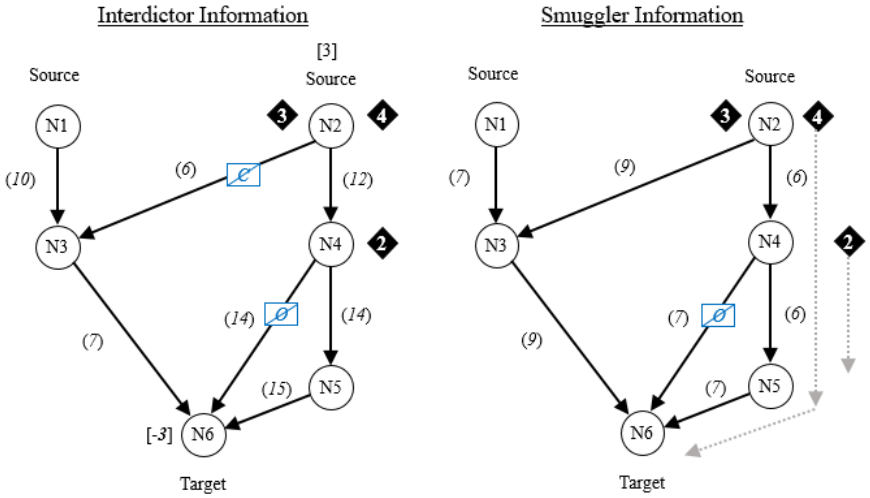

Left (Interdictor): Through a feedback loop, the interdictor quickly refines his estimate of arc attributes, and, in spite of partial information, places two sensors in a worst-case manner for the smuggler. *Right (Smuggler)*: In Case 2 the transformation function only updates the smuggler's private arc information. He attempts to move two packets to the target, but again only has sufficient budget to move one packet, packet 4, all the way to the target, recording a score of 1.

Figure 13. Case 2, Round 3 Decisions and Information.

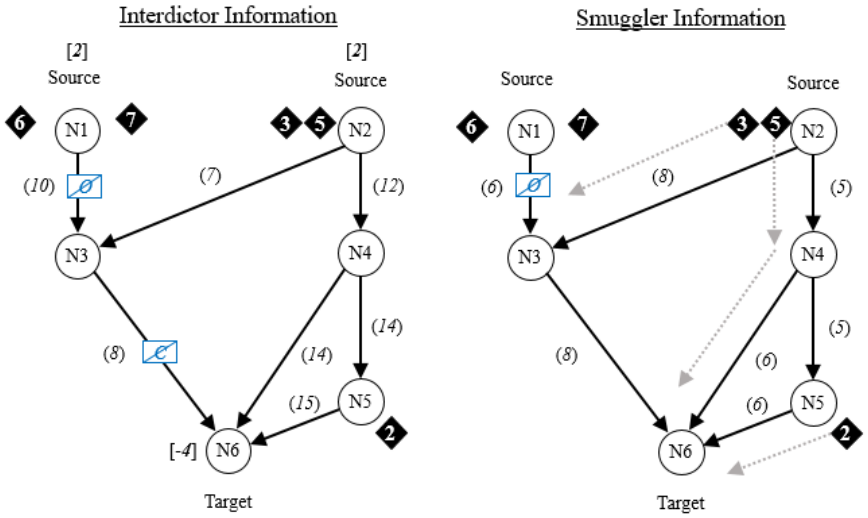

Left (Interdictor): The interdictor has updated his estimates again. However, in spite of perfect information of the packet locations, places sensors sub-optimally, as he did in round 1. *Right (Smuggler)*: The transformation function again updates the smuggler's information. Limited by budget, the smuggler moves two of three packets with seemingly unobstructed paths to the target, recording a score of 2.

Discussion

In Case 2, we assume the interdictor is aware of all packet locations. However, he must form his own estimate of the arc costs and capacities. The now-inaccurate arc information causes the interdictor to place a covert sensor inappropriately in Round 1 (Figure 11). Were the interdictor aware of the smuggler's arc information, he would have placed the covert sensor more appropriately on arc (*N2, N4*) as in Case 1, Round 1 (Figure 7). The consequences of the miscalculation allow the smuggler to take advantage of an unimpeded path by moving packet 1 to the target for score. Even so, the interdictor is able to rapidly adjust his estimate and more-appropriately place sensors in the next round. The introduction of asymmetric arc cost and capacity information immediately shows a difference in the accuracy of decisions made by the interdictor and smuggler. The model in Case 2 bears similarity to previous work using limited feedback (e.g., Bubeck 2011). We argue that a model of the interdictor-smuggler problem must include asymmetric information; however, assuming that the interdictor knows the location of all smuggled materiel as it moves is problematic. We examine this further in Case 3.

3. Case 3: Multi-period, Symmetric Cost and Capacity Information, Packet Locations Unknown to Interdictor, Limited Feedback

In Case 3, we assume the smuggler and interdictor have access to the same arc costs and capacities. However, packet locations are private, known only to the smuggler. The interdictor uses the size and location of destroyed packets as feedback to estimate the smuggler's supply nodes. The smuggler can observe overt sensor locations, but cannot see covert sensor locations. The smuggler uses the same cost and capacity feedback loop found in Case 2 to update his estimates of arc costs. The interdictor has access to these estimates of arc cost and capacity. The entire set of arcs is visible to the smuggler throughout the game. Figure 14 displays the time-ordered steps of the game algorithm for constructive Case 3. Figures 15–17 display the decisions, information, and outcomes of three rounds played under Case 3.

Figure 14. The **I-I/S** Game Algorithm, Case 3.

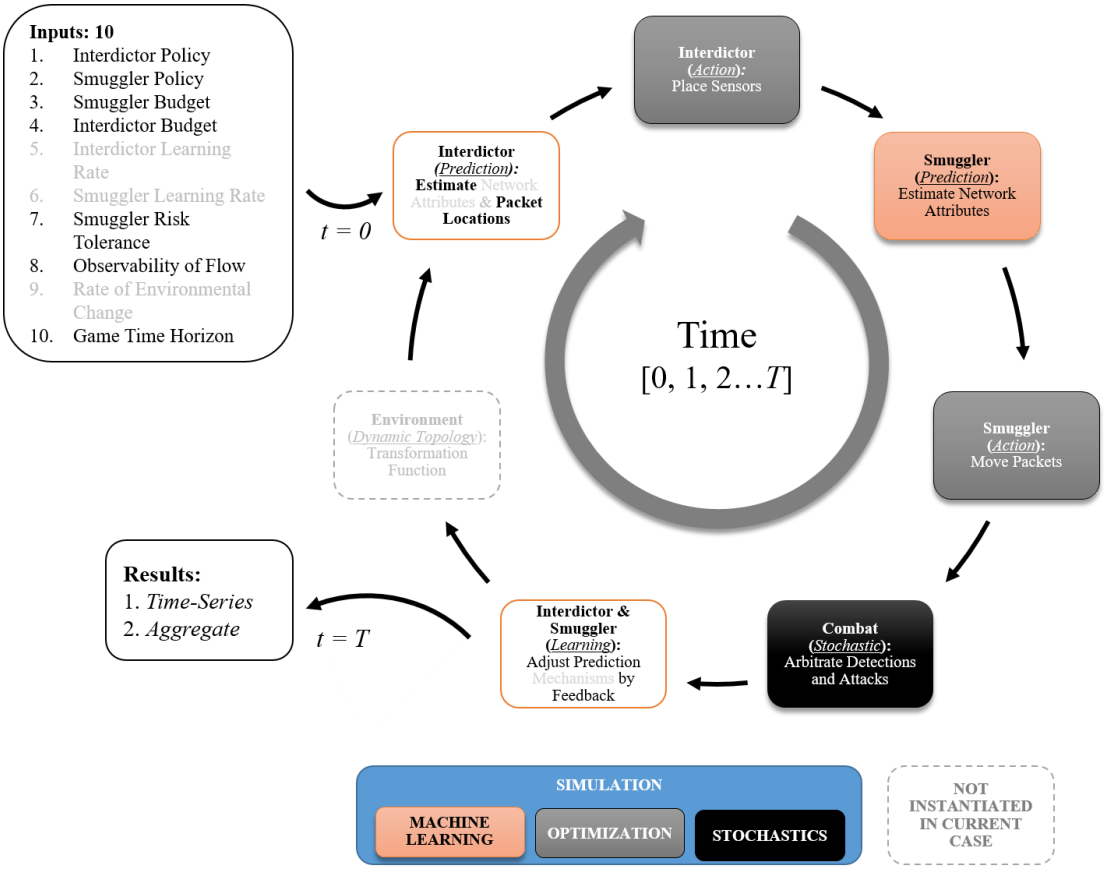

The time-ordered steps of the game algorithm for constructive Case 3. In Case 3, the information asymmetry is of a different type than that found in Case 2. The interdictor and smuggler share the same estimate of the arc attributes, but the location of the packets is private, known only to the smuggler. The interdictor uses a limited feedback mechanism to attempt to compensate for the incomplete information.

Figure 15. Case 3, Round 1 Decisions and Information.

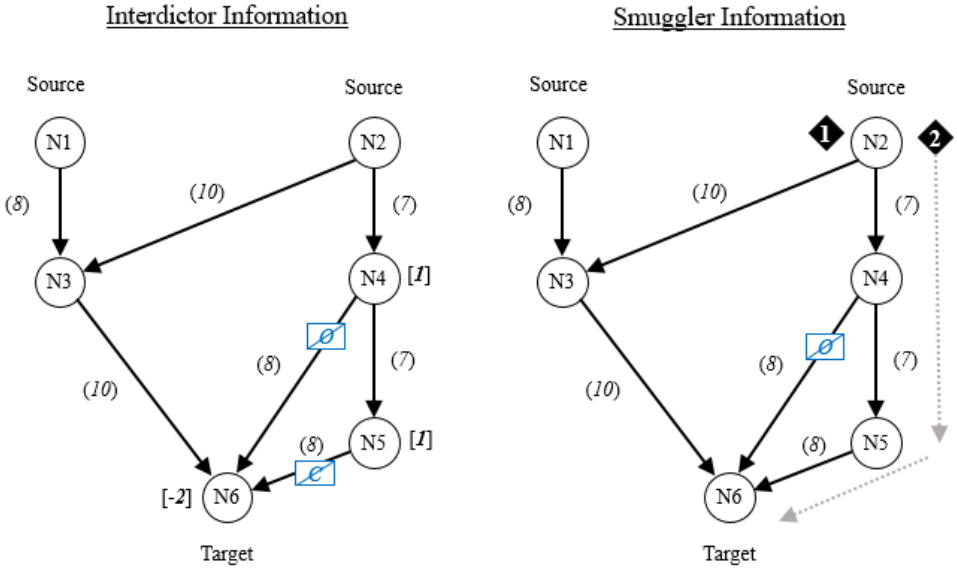

This case follows the same packet schedule as the previous ones. *Left (Interdictor)*: As in Case 2, the interdictor is aware of the smuggler's information on arc attributes. However, the interdictor can no longer see the packet locations and must estimate them. We suppose his initial guess is inaccurate, estimating 1 unit of supply each at nodes *N4* and *N5*. *Right (Smuggler)*: The smuggler is unable to see covert sensors. He attempts to move packet 2 to the target based on his partial information. Of note, the chosen *s*-*t* path (*N2*-*N4*-*N5*-*N6*), is not the shortest path. However, it moves the same amount of materiel forward as the shortest path (*N2*-*N3*-*N6*). We see some diversity among budget-feasible paths because the smuggler is attempting to move as much flow as far forward as possible within a budget that does not roll-over round-to-round. With a budget of 25 units, both paths (*N2*-*N3*-*N6*) and (*N2*-*N4*-*N5*-*N6*) fall within the smuggler's indifference threshold. He treats them equally. The stochastic result from packet 2 passing the covert sensor on arc (*N5*, *N6*) yields "no detection." Packet 2 reaches the target giving a score of 1.

Figure 16. Case 3, Round 2 Decisions and Information.

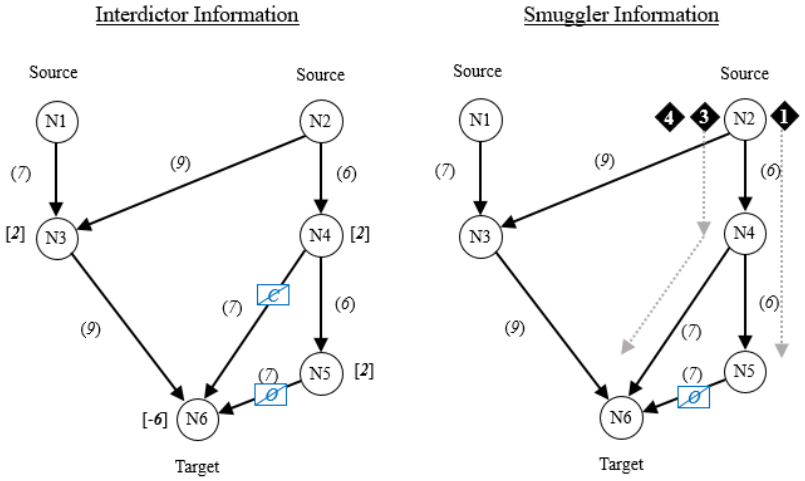

Left (Interdictor): The interdictor estimates two packets each at nodes *N3*, *N4*, and *N5*. However, without perfect knowledge of packet locations, the interdictor keeps sensors close to the target, sub-optimally placing them based on updated supply estimates at nodes *N4* and *N5*. *Right (Smuggler)*: The transformation function updates the smuggler's private estimate of arc attributes. He attempts to move packet 3 to the target, while moving packet 1 toward the target. The stochastic result from packet 3 passing the covert sensor on arc (*N4*, *N6*) yields "no detection." Packet 3 reaches the target, increasing the total score by 1.

Figure 17. Case 3, Round 3 Decisions and Information.

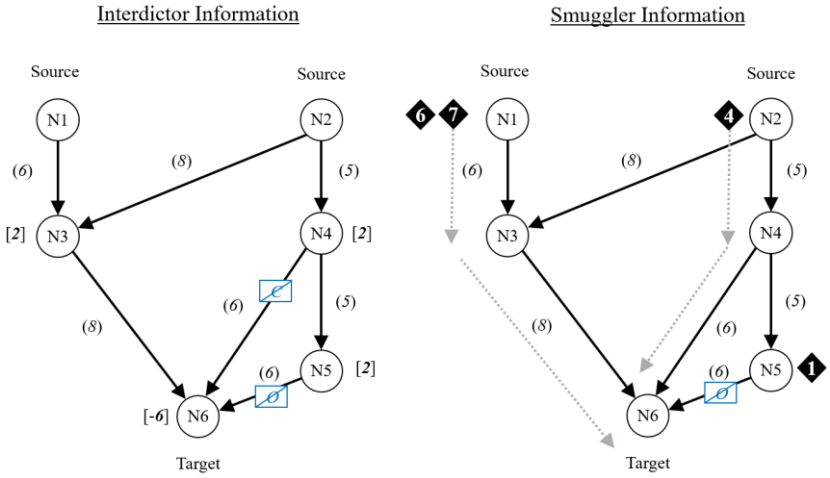

Left (Interdictor): Using feedback from round 2, the interdictor refines his estimate of packet locations. Because of persistent imperfections in this estimate, the interdictor again places sensors poorly, even with perfect information of the smuggler's arc attributes. *Right (Smuggler)*: After the transformation function updates the smuggler's information, the smuggler attempts to move two packets to the target. The stochastic result from packet 4 passing the covert sensor on arc (*N4*, *N6*) yields "no detection." Packets 4 and 7 reach the target, increasing the total score by 2.

Discussion

In Case 3, we assume the interdictor is now unaware of all packet locations, but sees the same arc costs and capacities as the smuggler. The now-hidden packets dramatically affect both the interdictor's and smuggler's decisions. Throughout Case 3, the interdictor places his sensors much closer to the target than in Cases 1 and 2. Uncertainty on the origin of smuggler packets greatly complicates the interdictor's problem because any of the five nodes could act as a source of flow. While the sensors are optimally placed given the interdictor's estimates of supply, their location is sub-optimal every round when compared to what could be achieved with the ground truth. Case 3 demonstrates the significant impact on both the decisions and outcome of the game when the interdictor has incomplete information of the packet's locations.

Additionally, Case 3 illustrates smuggler selection diversity amongst several budget-feasible paths. Given that the smuggler is unable to roll-over any excess movement budget round-to-round, two paths that move an equal amount of materiel forward in a game round both lie within the smuggler's indifference threshold. The smuggler may then equally choose either path (Figure 15). (For further information, see Stewart et al. [2013] for a detailed treatment of indifference thresholds and other multicriteria decision-making.)

The model in Case 3 is similar to previous models that assume the interdictor and smuggler both know the probability of detection on each arc but that the smuggler's origins are unknown to the interdictor (e.g., Morton et al. 2007). The obscurity of packet locations assumed in Case 3 aligns well with the author's own observations in combat during counter-trafficking and counter-infiltration operations. However, the assumption of perfect cost and capacity information is difficult to reconcile with the same operational experiences. We relax both assumptions in Case 4.

4. Case 4: Multi-Period, Asymmetric Cost and Capacity Information, Packet Locations Unknown to Interdictor, Limited Feedback

In Case 4, we combine the restrictions from Case 2 and Case 3. We assume the smuggler and interdictor make decisions with private arc cost and capacity information. They both use limited feedback from packets destroyed in each round to update their

information. The packet locations are not known to the interdictor and must be estimated. As in previous cases, the entire set of arcs is visible to the smuggler throughout the game. Figure 18 displays the time-ordered steps of the game algorithm for constructive Case 4. Figures 19–21 display the decisions, information, and outcomes of three rounds played under Case 4.

Figure 18. The **I-I/S** Game Algorithm, Case 4.

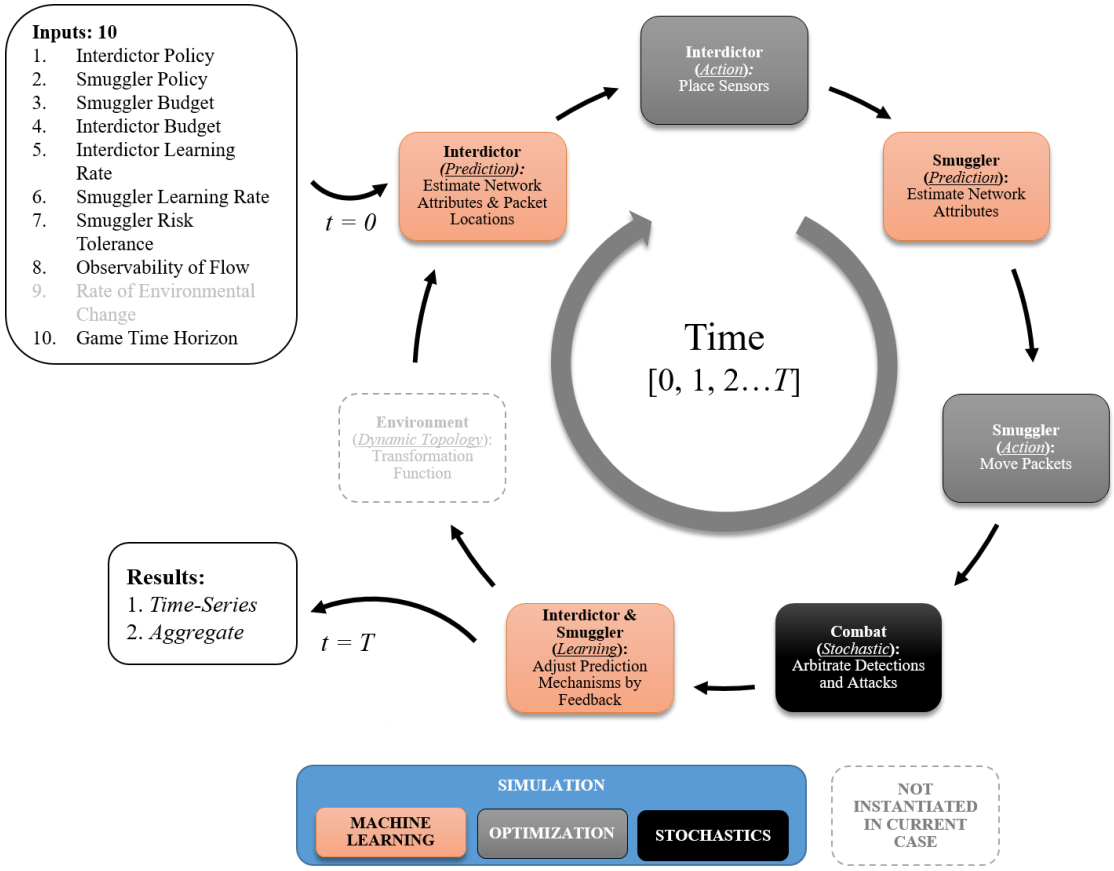

The time-ordered steps of the game algorithm for constructive Case 4. In Case 4, we combine the restrictions from Cases 2 and 3. The interdictor must now estimate both the arc attributes and packet locations. The interdictor uses two limited feedback mechanisms to update his projections of the smuggler's capabilities and intentions.

Figure 19. Case 4, Round 1 Decisions and Information.

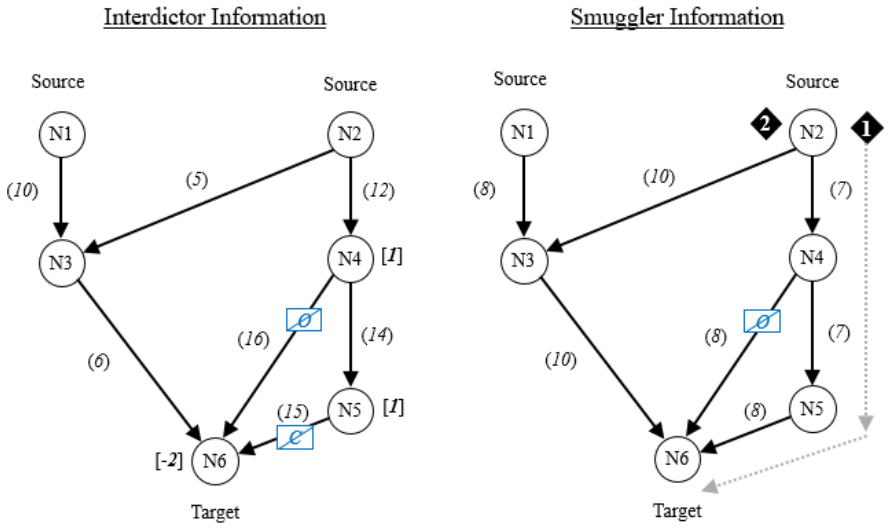

Left (Interdictor): The interdictor must estimate both the packet locations and arc attributes in Case 4. He places two sensors inaccurately because of the lack of information. *Right (Smuggler)*: The smuggler attempts to move packet 1 to target, without awareness to the covert sensor. The stochastic result from packet 1 passing the covert sensor on arc (*N5*, *N6*) yields "no detection." Packet 1 reaches the target, increasing the total score by 1.

Figure 20. Case 4, Round 2 Decisions and Information.

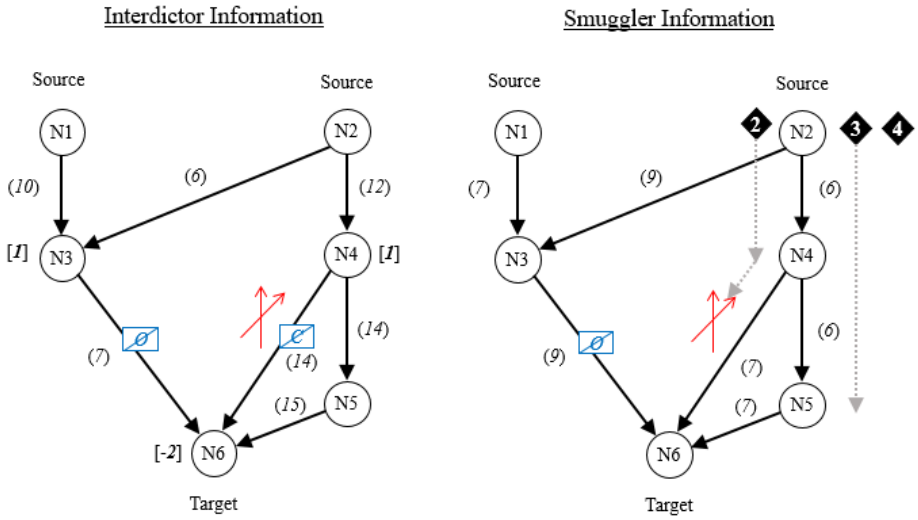

Left (Interdictor): Now employing two feedback mechanisms, the interdictor more accurately estimates the smuggler's intentions and well-places two sensors on arcs (*N3*, *N6*) and (*N4*, *N6*). *Right (Smuggler)*: The smuggler attempts to move two packets to target, without awareness of the covert sensor on arc (*N4*, *N6*). The stochastic result is that the interdictor detects and destroys packet 2 on arc (*N4*, *N6*).

Figure 21. Case 4, Round 3 Decisions and Information.

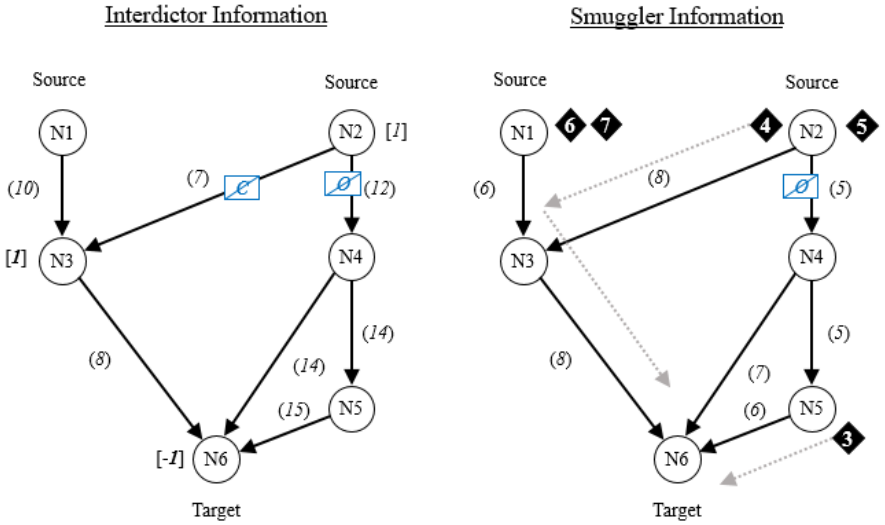

Left (Interdictor): Based on previous two rounds of play, the interdictor orients on node *N2* as the source of supply. He is surprised by flow originating instead from nodes *N1* and *N5* and fails to guard against it. *Right (Smuggler)*: The smuggler updates his information based on the packet lost in round 2 and the transformation function. He then attempts to move two packets to the target. The stochastic result is that the interdictor does not detect packet 2 on arc (*N2, N3*). Note that packets initially frustrated in the network, such as packet 3, may become a threat later. In this case, packet 3 reaches the target after the interdictor places sensors further forward.

Discussion

In Case 4, the total asymmetry of information again changes the decisions and outcomes during game play. Most notably, the interdictor benefits from imperfect information on arc cost and capacity. Without direct access to the smuggler's information, the interdictor must use an additional feedback mechanism to form an estimate of these parameters based on game play. One feedback mechanism is already in use to discern the location of the packets. When the two are combined, the interdictor is able to make better projections by using complimentary information. Estimated changes in cost resultant from observed packets enable the interdictor to refine the estimates of flow sources and lead to better-placed sensors. The interdictor makes decisions based on where the smuggler *did go*, not just where he *could go*. There were no models found during the literature review that included the level of information asymmetry discussed above.

5. Case 5 (Full Model): Multi-period, Asymmetric Cost and Capacity Information, Packet Locations Unknown to Interdictor, Limited Feedback, Arcs Revealed by Timer

Case 5, the full model, includes all of the attributes of Case 4. However, unlike Case 4, we assume a count-down timer makes some previously-hidden arcs visible to the smuggler as the game progresses. Once visible, arcs are always visible. Figure 22 displays the time-ordered steps of the game algorithm for constructive case 5. Figures 23–25 display the decisions, information, and outcomes of three rounds played under Case 5.

Figure 22. The **I-I/S** Game Algorithm, Case 5.

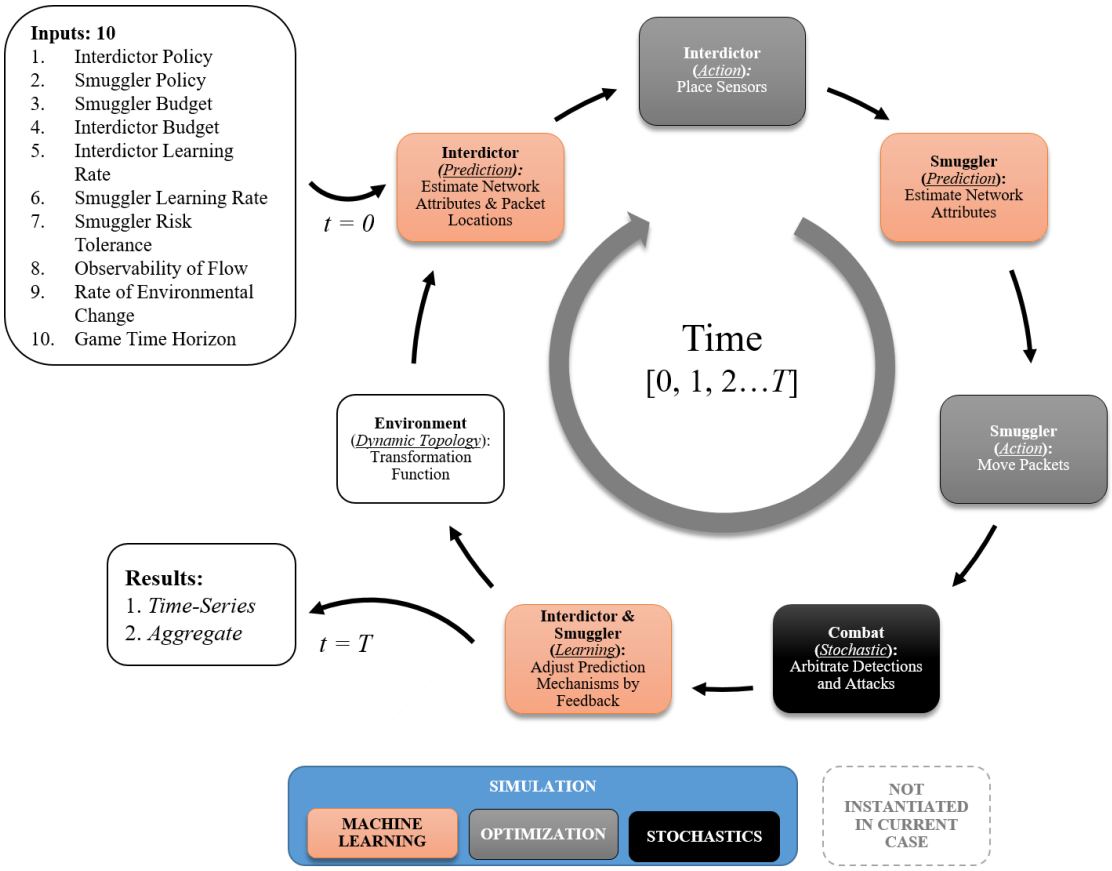

The time-ordered steps of the game algorithm for constructive Case 5 (the full model). In Case 5, we include all of the attributes of Case 4 but now use a countdown timer to hide some arcs from the smuggler until round 3.

Figure 23. Case 5, Round 1 Decisions and Information.

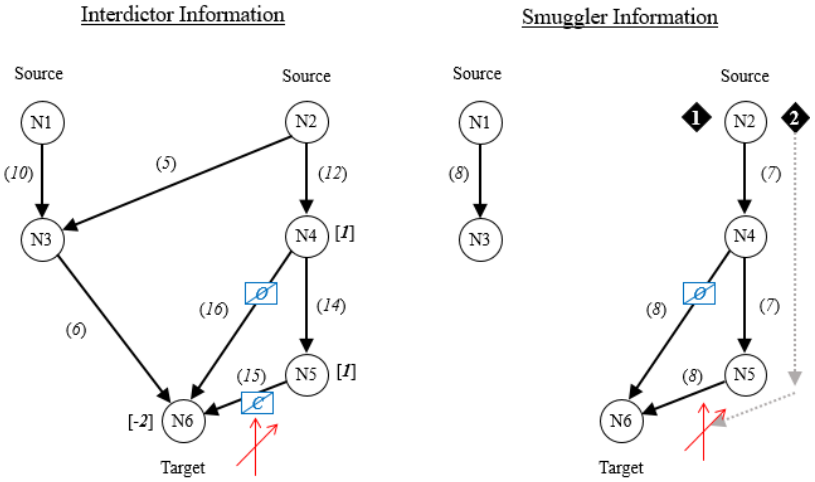

Left (Interdictor): As in Case 4, the interdictor must estimate both the packet locations and arc attributes. He places an overt sensor on *(N4, N6)* and covert sensor on *(N5, N6).* While these arcs are optimal sensor locations given the interdictor's current information and estimate, they are sub-optimal locations in view of the actual position of the two packets. *Right (Smuggler)*: In Case 5, a clock will reveal arcs (*N2, N3*) and (*N3, N6*) to the smuggler by countdown timer in round 3. These arcs are not yet visible in round 1. The smuggler is thus far more limited in available routes. By chance, the interdictor's initial estimate aligns well with the limited routes available to the smuggler. The stochastic result from packet 2 passing the covert sensor on arc *(N5, N6)* yields "detection." The interdictor destroys the detected packet.

Figure 24. Case 5, Round 2 Decisions and Information.

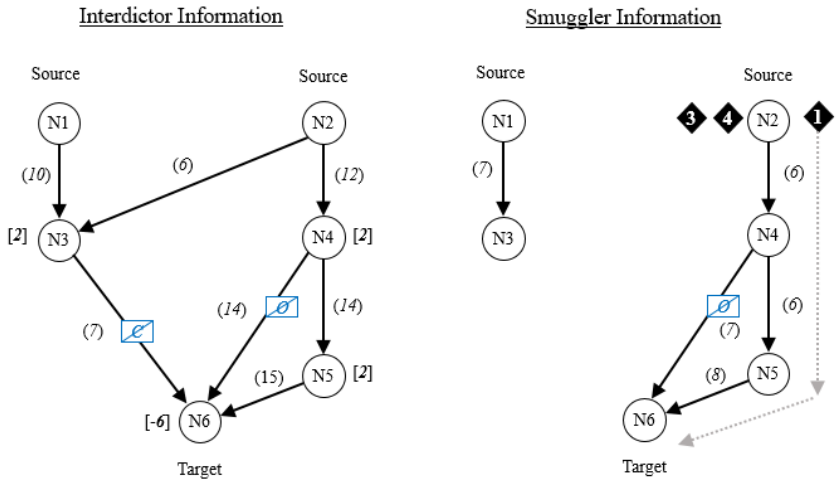

Left (Interdictor): With one detection in the previous round, the interdictor widens his estimate of supply nodes to include *N3*. The interdictor is unaware that arc (*N3, N6*) is not yet visible to the smuggler and treats it as a threat. The interdictor adjusts his estimates of arc capacities based on detected flow. The interdictor places two sensors using his current information. The placement is again sub-optimal, given the ground truth. *Right (Smuggler):* The smuggler successfully moves packet 1 to the target.

Figure 25. Case 5, Round 3 Decisions and Information.

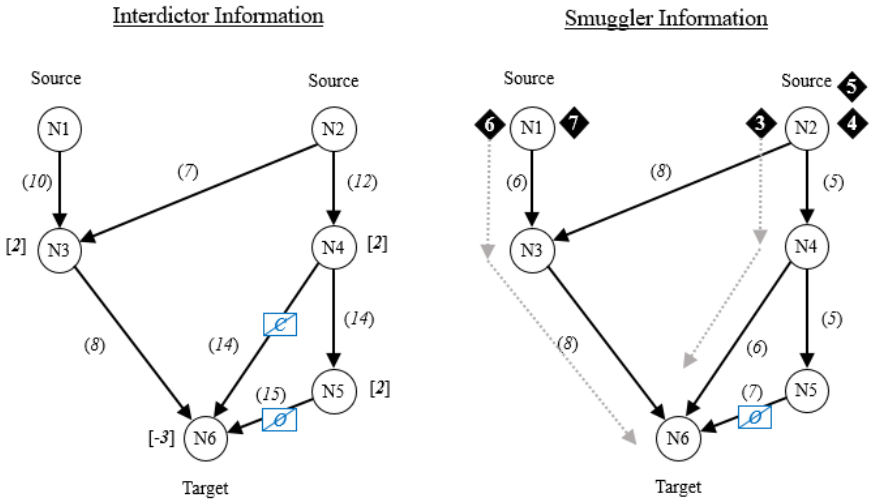

Left (Interdictor): Even with no detections in the previous round, the interdictor is able to use his estimate of arc attributes, arc supplies, and score from the previous round to improve the accuracy of his information. He correctly identifies node *N2* as a source of supply and nodes *N3* and *N5* as threatening. However, the limited smuggling routes that allowed the interdictor to gain a detection also created feedback that causes him to orient his defense without regard for node *N1* as a possible source. *Right (Smuggler):* The countdown timer reveals arcs (*N2, N3*) and (*N3, N6*) to the smuggler. The smuggler is able to capitalize on the uninhibited *s-t* path (*N1-N3-N6*) created by the interdictor's now-obsolete sensor orientation to move packet 6 to the target6. The stochastic result of packet 3 passing the covert sensor on arc (*N4-N6*) is "no detection."

Discussion

The time-revelation of new arcs in Case 5 again changes game play. The introduction of previously unknown arcs is analogous to opening a new cross-border smuggling tunnel. There are other real world examples of new smuggling routes coming into use (e.g., Banco 2015). The addition of arcs revealed by timer most affects the smuggler's decisions. By constricting the available set of arcs, the smuggler becomes more predictable. In time, the increased predictability would make the interdictor more successful. However, once the interdictor is oriented on a specific set of arcs, the smuggler's use of new routes could prove a surprise to the interdictor's established sensor layout. The inclusion of time-delayed arc visibility allows for the investigation of this phenomenon and aligns with real world events.

Based on discussion in the above constructive examples, we argue that the assumptions found in Case 5, the full model, are each significant and well represent the key features underlying the interdictor-smuggler problem we describe.

D. MATHEMATICAL FORMULATION OF THE I-I/S ALGORITHM

In this section, we describe the detailed mathematical formulation of each step of the **I-I/S** Simulation Algorithm (Figure 26).

Figure 26. The Steps of the **I-I/S** Simulation Algorithm.

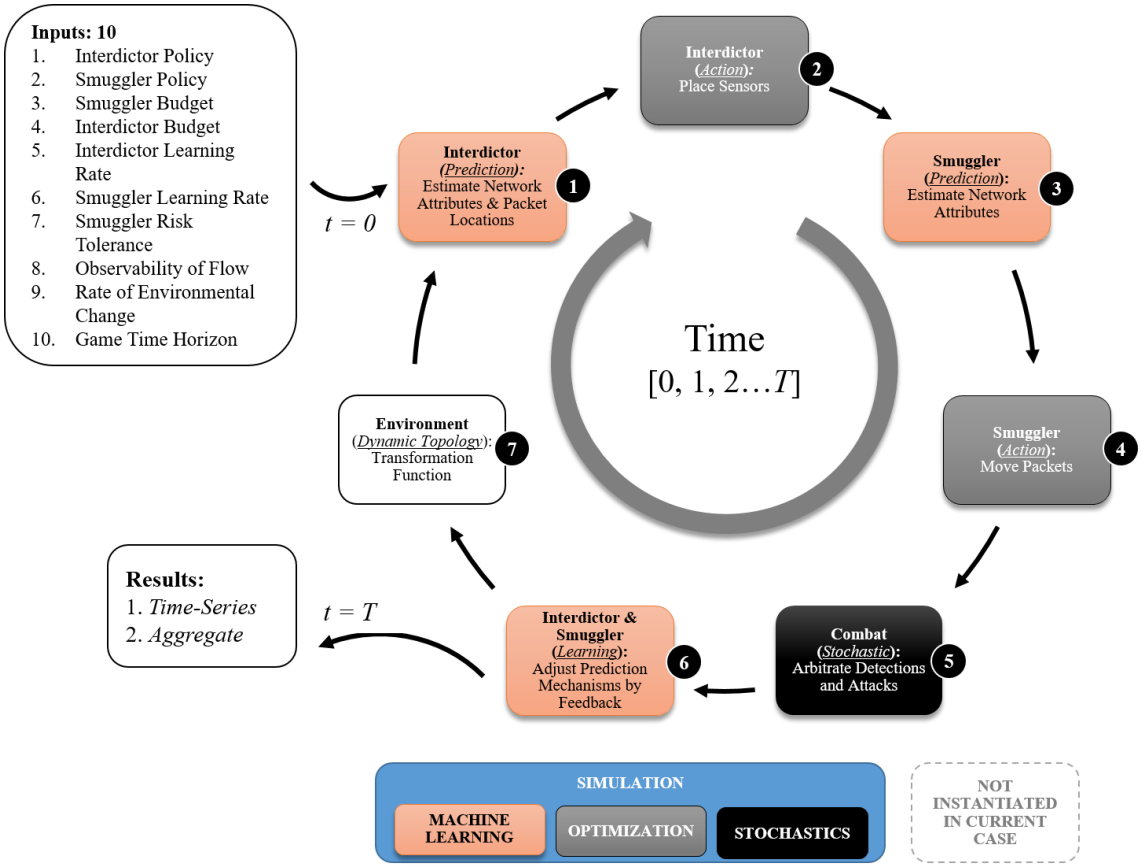

The numbered steps of the simulation algorithm. Each number corresponds with the subsection describing its detailed mathematical formulation.

1. (Interdictor) Estimate Network Attributes and Packet Locations

a. Introduce New Flow to the Source(s)

A predefined *master schedule* controls flow input. We define the master schedule a priori and all basic packet attributes within it. It is in table form. Table 3 provides an illustrative sample. The smuggler has no foreknowledge of the master schedule.

Table 3. Illustrative Sample of Master Packet Flow.

Round	Packet	Source	Target	Size
1	P1	N1	N6	1
1	P2	N1	N6	2
2	P3	N2	N6	3
3	P4	N1	N6	3

All packet attributes are assigned in the master packet flow table.

In the above example, there are two sources of flow, nodes *N1* and *N2*, and one target, node *N6*. The algorithm will place two packets, *P1* and *P2*, in the network on round 1 at node *N1*. These packets are sizes 1 and 2, respectively. The algorithm will add one packet to the network in each rounds 2 and 3. *P3* will be added at node *N2* and packet *P4* will be added at node *N1*. Once the flow has been added to the network, it is no longer directly influenced by the master schedule. Lastly, the algorithm assigns both the interdictor and smuggler their round-to-round budgets in Step 1.

b. (Interdictor) Estimate Node Supplies (Algorithm ESTIMATE_SUPPLIES)

The following sub-algorithm calculates the interdictor's estimated node supplies and demands:

Algorithm ESTIMATE_SUPPLIES

(1) Solve Balance of Flow Equation for Supply

Sets and Indices

$i \in N$ node (alias *j*, *nodes*)

$(i, j) \in A$ arc directed from node *i* to node *j*

$t \in T$ game round in ordered set of total rounds, *T*. Under control of algorithm **I-I/S**

Data [units]

$flow_{(i,j),t-1}$ flow detected in previous round on arc (*i, j*) [flow/round]

Formulation

$$supply_{i,t} = \sum_{(i,j)\in A} flow_{i,j,t-1} - \sum_{(i,j)\in A} flow_{j,i,t-1} \quad (1.1)$$

(2) If Sensor Emplaced on Arc, Update Network Supply Estimates

For node *i* in reverse adjacency list of node with sensor in forward star (*FS*)

$$supply_{i,t} = \max(supply_{j,t-1}) \quad j \in N \mid j \ \ adj \ \ i \quad (1.2)$$

(3) Remove Old Supply Estimate.

For node $i \in G(N)$

If $\exists$ sensor in $FS_i \wedge \nexists$ sensor in Reverse Star$_i$ (*RS$_i$*) **Then**

$$supply_{i,t} = 0 \quad (1.3)$$

Discussion

The algorithm calculates the interdictor's estimate of each node's supplies and demands from the total materiel sensed along each arc in the previous round (1.1). This is a simple balance of flow equation where the left hand side, $supply_{i,t}$, is the only unknown. Using these supplies and demands, the algorithm now extrapolates the interdictor's estimate to potential "upstream" flow origins. In the **I-I/S** model, we assume that the interdictor is aware of arcs in the reverse star of any node with a sensor in the forward star (1.2) (Figure 27).

Figure 27. Arcs Visible to the Interdictor.

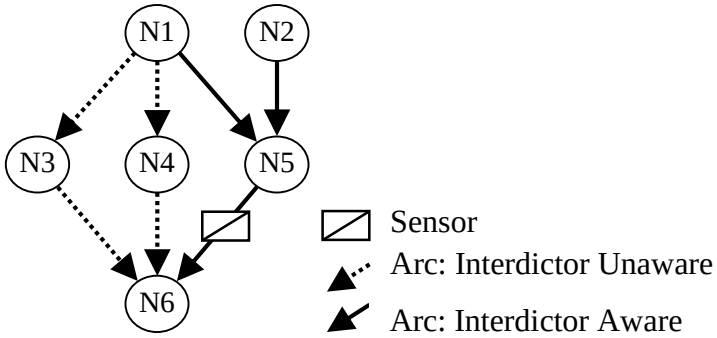

A sensor placed on arc (*N5, N6*) makes the interdictor aware of the arcs in the reverse star of node *N5*, arcs (*N1, N5*) and (*N2, N5*). The interdictor is unaware of the remaining arcs.

Lastly, the algorithm reconciles the interdictor's new estimates of supplies and demands to avoid duplication of these supplies and demands among neighbors (1.3). Figure 28 illustrates algorithm ESTIMATE_SUPPLIES by example.

Figure 28. Example of Algorithm ESTIMATE_SUPPLIES.

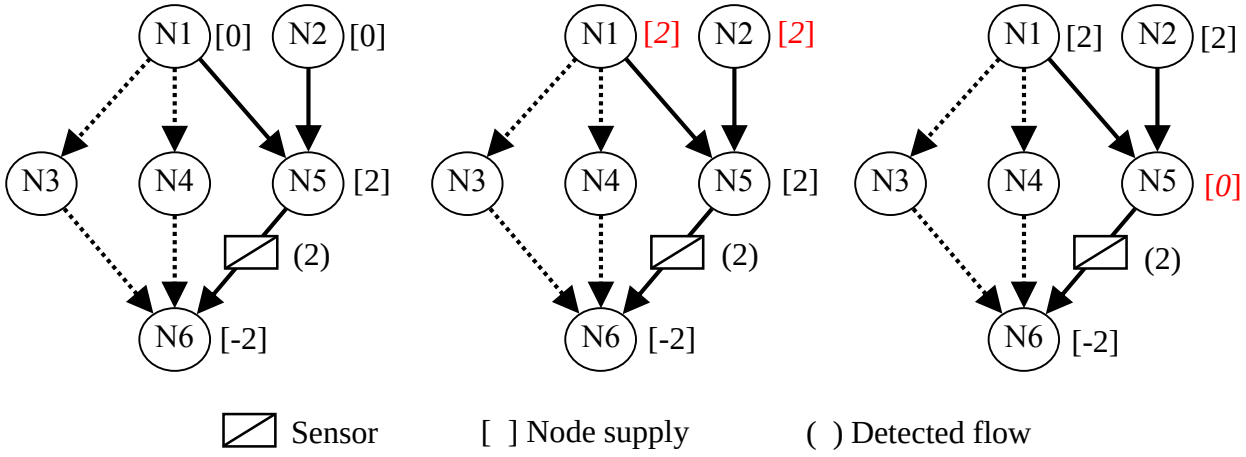

Algorithm ESTIMATE_SUPPLIES proceeds left to right. Values that change are highlighted in *red italics*. ***Step a***: At left, the algorithm calculates node *N5* and *N6* supplies as 2 and -2 based on 2 units of flow being observed on arc (*N5, N6*). ***Step b***: Center, the algorithm uses this estimate to extrapolate supplies at nodes *N1* and *N2*. ***Step c***: Right, the algorithm reconciles all node supplies, setting node *N5* supply to 0.

2. (Interdictor) Place Sensors

In Step 2, the interdictor places both overt and then covert sensors subject to a sensor budget.

a. (Interdictor) Set Covert Sensors Subject to Budget

The interdictor places covert sensors optimally, given his available partial information, by solving a minimum cost flow interdiction problem as below. The minimum cost flow interdiction sub-problem places sensors upon arcs instead of performing attacks against them as in traditional minimum cost flow interdiction formulations. Solving the SENSOR_PLACEMENT Dual Integer-Linear Program (ILP) yields the sensor placement plan with the maximum estimated minimum-cost of flow. The master **I-I/S** algorithm controls round indices, $t \in T$.

Sets and Indices

$i \in N$	node (alias *j, nodes*)
$(i, j) \in A$	arc directed from node *i* to node *j*
$t \in T$	game round in ordered set of total rounds, *T*. Under control of algorithm **I-I/S**

Data [units]

$\hat{c}_{(i,j),t}$	interdictor's estimate of smuggler's cost to move packet of size 1 across arc (i, j) in round *t* [cost/flow-round]
$\hat{u}_{(i,j),t}$	interdictor's estimate of smuggler's capacity on arc (i, j) in round *t* [flow/round]
$\hat{b}_{i,t}$	interdictor's estimate of supply at node *i*, >0 supply at node *i*, <0 demand at node *i*, = 0 transshipment [flow/round]
penalty	penalty for flow across arc (*i*, *BFN*) or (*BFN*, *j*) | $i, j \in N$
$covert.budget_t$	maximum number of covert sensors the interdictor can place in round *t* [cardinality/round]
$q_{(i,j),t}$	penalty for traversing an arc with sensor [cost/flow-round]

Decision Variables [units]

$X_{(i,j),t}$ flow from node *i* to node *j* in round *t* [flow/round]

$A_{(i,BFN),t}$ flow from node *i* to "Big Fake Node" (*BFN*) in round *t* [flow/round]

$E_{(BFN,j),t}$ flow from "Big Fake Node" to node *j* in round *t* [flow/round]

$Covert.Sensor_{i,j,t}$ =1 if place sensor on arc $(i, j) \in A$ during round *t*, = 0 otherwise [binary]

Maximin Formulation to increase arc flow costs [Dual Variables]

$$\max_{Y \in \Upsilon} \left\{ \begin{array}{l} \min\limits_{(i,j)\in A} \left[\left(\hat{c}_{i,j,t} + q_{i,j,t} Covert.Sensor_{i,j,t} \right) X_{i,j,t} + \sum\limits_{i \in N} penalty \cdot (E_{i,BFN,t} + A_{BFN,i,t}) \right] \\ s.t. \quad \sum\limits_{(i,j)\in A} X_{i,j,t} - \sum\limits_{(j,i)\in A} X_{j,i,t} + E_{i,BFN,t} - A_{BFN,i,t} = \hat{b}_{i,t} \quad \forall i \in N \qquad [\pi_{i,t}] \\ \qquad 0 \le X_{i,j,t} \le \hat{u}_{i,j,t} \qquad \forall (i,j) \in A \quad [\alpha_{i,t}] \\ \qquad E_{i,BFN,t}, A_{BFN,i,t} \ge 0 \end{array} \right\} \quad (1.4)$$

$$\Upsilon = \left\{ \begin{array}{l} \sum\limits_{(i,j)\in A} Covert.Sensor_{i,j,t} \le covert.budget_t \\ \qquad Covert.Sensor_{i,j,t} \in \{0,1\} \quad \forall (i,j) \in A \end{array} \right\} \quad (1.5)$$

SENSOR_PLACEMENT Dual ILP Formulation [Primal Variables]

$$\max_{\alpha,\beta,Y} \quad \sum_{i \in N} \hat{b}_{i,t} \pi_{i,t} - \sum_{(i,j)\in A} \hat{u}_{i,j,t} \alpha_{i,j,t} \quad (1.6)$$

$$s.t. \quad \pi_{i,t} - \pi_{j,t} - \alpha_{i,j,t} - q_{i,j,t} Covert.Sensor_{i,j,t} \le \hat{c}_{i,j,t} \quad \forall (i,j) \in A \quad [X_{i,j,t}] \quad (1.7)$$

$$\sum_{(i,j)\in A} Covert.Sensor_{i,j,t} \le covert.budget_t \quad (1.8)$$

$$\alpha_{i,j,t} \ge 0 \qquad \forall (i,j) \in A \quad (1.9)$$

$$Covert.Sensor_{i,j,t} \in \{0,1\} \qquad \forall (i,j) \in A \quad (1.10)$$

$$-penalty \le \pi_{i,t} \le penalty \qquad \forall i \in N \quad (1.11)$$

Discussion

We base the interdictor's maximin formulation on the interdictor's estimated supplies and demands, $\hat{b}_{i,t}$, costs, $\hat{c}_{i,j,t}$, and capacities, $\hat{u}_{i,j,t}$ (1.4). Using these estimated parameters, the interdictor places covert sensors to maximally penalize various potential interdictor flows subject to a budget (1.5). Examining the dual formulation reveals an ILP that is more easily solved optimally. We leverage elastic programming through variables $E_{i,BFN}$ and $A_{BFN,i}$. Each represents flow through an artificial sink node, *BFN*. The cost of this flow along arcs $(BFN, i) \mid i \in N$ is $n \cdot C_t$, where *n* is the order of the network and C_t the maximum of all arc costs during round *t*. This bounds the dual variables ensuring feasibility within the primal problem. The master **I-I/S** algorithm utilizes the resultant solution.

b. (Interdictor) Set Overt Sensors Subject to Budget

The interdictor reduces the penalty on arcs with covert sensors set from ***Step 2.a*** (above) to zero. This prevents the interdictor from selecting the same arc for both covert and overt sensors. The interdictor then finds overt sensor locations by solving another minimum cost flow interdiction problem with updated arc penalties.

Algorithm SET_OVERT_SENSORS

For $(i, j) \in A$

If $Covert.Sensor_{i,j,t} = 1$ **Then**

$q_{(i,j)} = 0$

$covert.budget_t = overt.budget_t$

Solve SENSOR_PLACEMENT Dual ILP

Discussion

The SENSOR_PLACEMENT Dual ILP formulation is the same, except, $Covert.Sensor_{i,t}$ is replaced by $Overt.Sensor_{i,t}$ in (1.4), (1.5), (1.7), (1.8), and (1.10); and

$covert.budget_{i,t}$ is replaced by $overt.budget_{i,t}$ in (1.8). As in ***Step 2.a*** of the **I-I/S** master algorithm above, these binary variables indicate the placement of overt sensors on arc (*i*, *j*) in round *t* if $Overt.Sensor_{i,j,t} = 1$. Similarly, the parameter $overt.budget_{i,t}$ limits the number of available overt sensors in round *t*.

3. (Smuggler) Estimate Network Attributes

A modified depth-first search (DFS) algorithm computes all simple paths from each current packet location to the packet's assigned target node. The master **I-I/S** algorithm uses the resultant *path* as an input to Step 4. Next, the smuggler increases his arc cost estimates for those arcs with overt sensors placed upon them. Covert sensors are not visible to the smuggler so the smuggler does not adjust his estimated arc costs in this case.

a. (Smuggler) For each Packet in Play, Compute all Simple Paths to the Target Node

Data [units]

$location_{p,i,t}$ =1 if packet *p* located at node *i* during round *t*, = 0 otherwise [binary]

Algorithm PACKET_PATHS

For $p \in Packets$

$k = i \mid (location_{p,i} = 1)$

list ={ }; *paths* = { }

SimplePaths (*i*, target, *list*, *paths*)

list = *list* $\cup$ *i*

If *i* == *"target"* **Then**

paths = *paths* $\cup$ *list*

Else

For *j* adjacent to *i* and $\notin$ *paths*

Call SimplePaths (*j*, target, *list*, *paths*)

list = *list* \ { *i* }

b. Increase Smuggler Arc Costs.

For $(i, j) \in A$

If $Overt.Sensor_{i,t} = 1$ **Then**

$$c_{(i,j),t} = c_{(i,j),t-1} + q_{(i,j)} \tag{1.12}$$

Discussion

The simple recursive algorithm above continues a DFS from the packet's current location, node *i*, through each node *j* adjacent to node *i* until reaching the target. Next, the algorithm examines candidate paths and retains them if they are unique and lead to the target, discarding all other paths. The remaining data container *paths* enumerate the set of simple paths. The PACKET_PATHS algorithm begins by increasing the smuggler's cost over arcs upon which the interdictor has placed an overt sensor (1.12). This provides awareness of these emplacements to the smuggler.

4. (Smuggler) Move Packets

Next, the smuggler moves packets optimally in the network given his incomplete information. The smuggler determines these movements by solving the following optimization problem. Note that within the objective function, material moved *to the target* is rewarded more so than that moved only *toward the target*.

Solve PACKET_MOVES

Sets and Indices

$i \in N$	node (alias *j*, *nodes*)
$(i, j) \in A$	arc directed from node *i* to node *j*
$p \in Packets$	packet
$short_{p,(i,j),t} \in Short.paths_t$	arc in simple path for packet *p* from current location to target
$touchdown_{(i,target)} \subset A$	arc in reverse star of target node

$t \in T$	game round in ordered set of total rounds, *T*. Under control of algorithm **I-I/S**

Data [units]

$size_p$	size of packet *p* [flow]
$location_{p,i,t}$	=1 if packet *p* located at node *i* during round *t*, = 0 otherwise [binary]
$c_{(i,j),t}$	smuggler's cost to move packet of size 1 across arc (i, j) during round *t* [cost/flow-round]
$u_{(i,j),t}$	smuggler's capacity on arc (i, j) during round *t* [flow/round]
$budget_t$	smuggler's *movement budget* during round *t*

Decision Variables [units]

$Move_{p,(i,j),t}$	=1 if move packet *p* on arc (*i, j*), =0 otherwise [binary]

PACKET_MOVES Formulation

$$\min_{Move} \sum_{p,(i,j)|p,(i,j)\in Short.paths} [size_p \cdot (1 - Move_{p,(i,j),t}) \cdot c_{(i,j),t} - budget_t \cdot Move_{p,(i,j),t|(i,j)\in touchdown} \cdot size_p] \quad (1.13)$$

$$s.t. \quad Move_{p,(i,j),t} - location_{p,i,t} \le 0 \qquad \forall p,(i,j) \in location_{p,i,t} \quad (1.14)$$

$$\sum_{nodes|(j,nodes)\in A} Move_{p,(j,nodes),t} - Move_{p,(i,j),t} \le 0; \qquad \forall p,i,j \in short_{p,(i,j),t} \quad (1.15)$$

$$\sum_{p,(i,j)|(i,j)\in A} size_p \cdot Move_{p,(i,j),t} \cdot c_{(i,j),t} \le budget_t \quad (1.16)$$

$$\sum_p size_p \cdot Move_{p,(i,j),t} \le u_{(i,j),t} \qquad \forall (i,j) \in A \quad (1.17)$$

$$\sum_j Move_{p,(i,j),t} \le 1 \qquad \forall p,i \quad (1.18)$$

Discussion

The first term of the objective function, $size_p \cdot (1 - Move_{p,(i,j),t}) \cdot c_{(i,j),t}$, calculates the smuggler's "work remaining" for each packet to reach its target (1.13). By summing over the all packets along the allowable paths, the cumulative work remaining is calculated. The second term, $-budget_t \cdot Move_{p,(i,j),t|(i,j)\in touchdown} \cdot size_p$, provides extra reward for the interdictor moving any packets fully to their assigned target. Binary decision variables with value 1, $Move_{p,(i,j),t} = 1$, reduce the amount of total remaining work and achieve the additional incentive of moving a packet to its target.

The first constraint ensures each packet begins movement from its present location (1.14). The second constraint mandates packet movement across consecutive arcs (1.15). Next, total packet movements are restricted by the smuggler's movement budget (1.16) and each arc's capacity (1.17). The last constraint prevents packets from taking multiple paths within any solution (1.18).

As illustrated in the constructive cases, the above formulation does not always cause the smuggler to utilize the lowest-cost paths. That is not the smuggler's objective. Several budget feasible paths yielding the same work remaining will all lie within the smuggler's indifference threshold. The smuggler is equally likely to choose any of these paths, creating path diversity in some instances. Path diversity within multicriteria decision-making problems well-aligns with the author's experience in tactical situations.

5. (Combat) Arbitrate Detections and Attacks

In Step 5, the **I-I/S** algorithm arbitrates any interactions between the smuggler's packets and the interdictor's sensors.

a. For those Packets that Moved Across a Sensor, Arbitrate Detections

First, the interdictor updates his current estimated cost for each arc $(i, j) \in A$. Then, for each packet, the algorithm evaluates the path assigned from Step 4 and sensor placements from Step 2. If there is a sensor on the path, the algorithm draws an outcome from the unique distribution describing the probability of detection for the packet on the

arc. A Tausworthe generator sets pseudo-random seeds (Tausworthe 1965). Begun with the aforementioned seeds, the algorithm draws pseudo-random variates from a Uniform distribution using a Mersenne Twister algorithm (Matsumoto 1998).

(1) Arbitrate Detections

$\hat{c}_{(i,j),t}$:	interdictor's estimate of arc (i, j) cost during round t [cost/flow-round]
$detections_{(i,j),t}$:	total amount of material detected on arc (i, j) during round t [non-negative integer]
packet.signature: $x \sim Uniform[0,1]$	level of stealth for packets [non-negative integer]
$size_{p,t}$:	size of packet p during round t [positive integer]

For packet p in current packets

For arc (i, j) in simple path assigned to packet p

If $Sensor_{(i,j),t} = 1$ **Then**

If $x \leq P^{detect}_{p,(i,j),t} = e^{0.1\left(-\hat{c}_{(i,j),t} + packet.signature \times size_{p,t}\right)}$ **Then** (1.19)

Packet p is detected on arc (i, j) in round t

(2) Tally Game Round Detections

For $(i, j) \in A$

$$detections_{(i,j),t} = \sum_{p | p \in detected.packets_{(i,j),t}} size_p \quad (1.20)$$

Discussion

The distribution describing the interdictor's probability of detection is a function of several variables. Equation (1.19) shows the probability of detecting a particular packet p in a specific round t on arc (i, j). As above, this probability is individually computed for each packet that is a candidate for detection (Figure 21). As described in Section A, different packets can have different sizes, but the size of each packet is fixed.

The parameter, $size_p$ thus allows the probability of detection to increase with increasing packet size. The parameter, *packet.signature* can further magnify the effect of size on the probability of detection. The magnification is important to isolate the influence that the ease of packet detectability plays in the interdictor-smuggler problem. The interdictor's estimate of arc (*i*, *j*) cost, $\hat{c}_{(i,j),t}$, changes round-to-round as a result of the number of previous detections. We detail the method of estimating cost change in Step 6. We mandate that $packet.signature \times size_p < \hat{c}_{(i,j),t}$ in equation (1.19) to ensure $P^{detect}_{p,(i,j),t} < 1$.

Figure 29. Sample Probabilities of Detection.

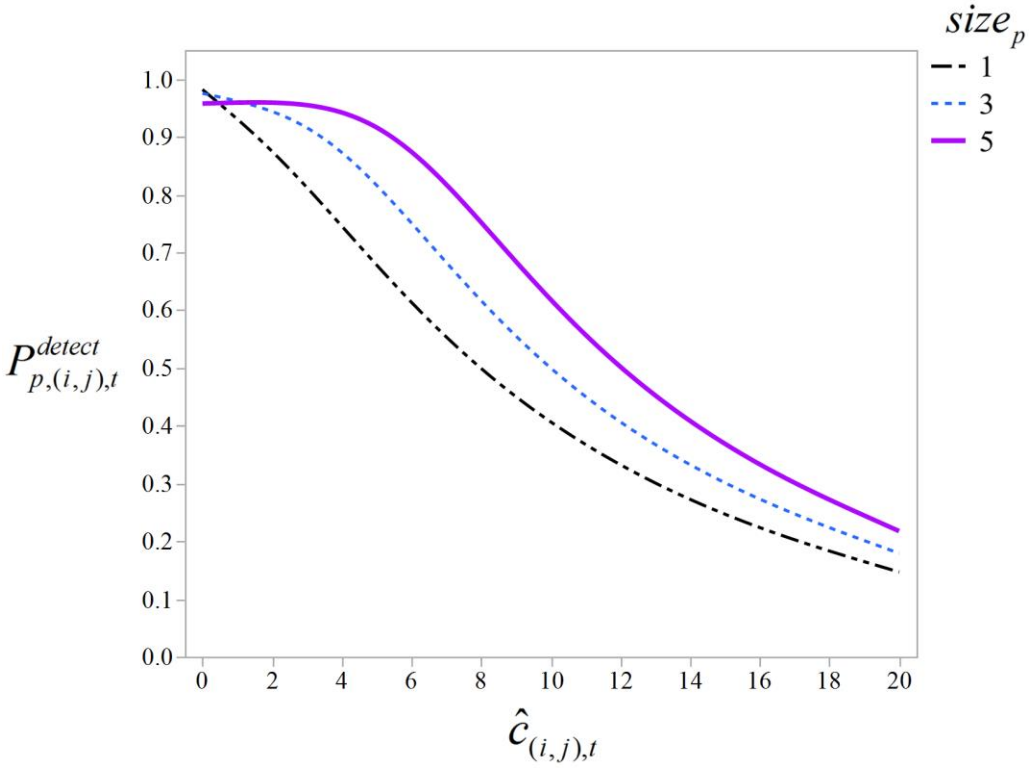

The plot above displays sample probabilities of detection for packets size 1, 3, and 5 by interdictor cost estimates. As size increases, the probability of detection increases. As estimated cost decreases, the probability of detection also increases.

b. (Interdictor) Attack and Remove Detected Packets

The interdictor "attacks" all detected packets. These attacks destroy the material within the packet, removing it from the game. Both the interdictor and smuggler tally the attacked materiel, affecting each of their estimated network information in Step 6. Attacks occur free of budget and are always successful.

For packet *p*

If packet *p* has been detected **Then**

Remove packet *p* from network

c. Tally Packets that Reach the Target Node and Remove Them from the Network

The algorithm removes packets from the network that have reached the target without destruction in the current round. The sum of these packet sizes, the *total materiel*, is the round's score. This ground truth score is available only to the smuggler. Feedback is limited, so the interdictor makes an estimate of the score in Step 6.

$$total.flow_t = \sum_{p|location_p = target} size_p \qquad \text{[flow/round]} \tag{1.21}$$

6. (Smuggler/Interdictor) Adjust Prediction Mechanisms by Feedback

The interdictor and smuggler update their private network information based on the results of play. First, the interdictor performs one final estimation of flow leaving the game through the target. Reconciliation of the remaining estimated network flow prevents a "death spiral" where the interdictor places fewer sensors by round and thus detects less flow by round until the flow estimate is zero and no sensors are placed. This occurs in spite of the actual non-zero flow reaching the target. Next, the smuggler's estimates of cost and capacity are influenced by both materiel lost and materiel successfully moved. Finally, the interdictor adjusts his estimated arc costs and capacities, completing the game round.

a. Calculate Interdictor's Estimate of Hits

$p.detected \subseteq Packets$: set of packets detected when reaching target and exiting the network in round t [cardinality]

$\widehat{total.hits_t}$: interdictor's estimate of material successfully reaching target in round t [flow/round]. $total.hits \leq total.flow_t$

$x \sim Uniform[0,1]$

For packet p:

If $location_{p,i,t} = target$ **Then**

If $x \leq P^{detect}_{p,target,t} = 1 - e^{-0.1\left(10+size_p\right)}$ **Then** (1.22)

Packet p is detected on exiting game through target in round t

$$\widehat{total.hits_t} = \sum_{p.detected|location_{p.detected}=target} size_p \quad (1.23)$$

Discussion

Equation (1.22) expresses the interdictor's probability of detecting packets that successfully hit the target and *exit* the game. As with the sensors' probability of detection, the probability of detection upon exit is a function of the packet's size; however, there is no direct effect of the detection history in this case (Figure 30). The uncertainty is necessary to mirror realistic smuggling scenarios. In this context an interdictor rarely has perfect information about the smuggler at any point—even at scenario's conclusion. Additionally, it treats situations where the interdictor is part of a layered defense and provided an estimate of materiel flow from the next line of defense. It requires the interdictor to learn and adapt under uncertainty, also facilitating miscalculation. These are important model facets to maintain realism.

Figure 30. Interdictor's Probability of Detecting a Packet Reaching the Target.

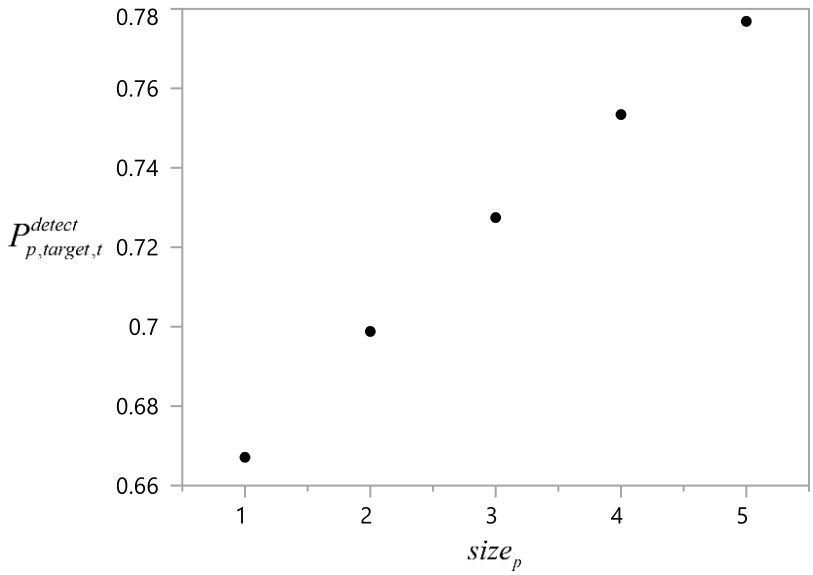

The plot represents only the probability of detecting packets that have otherwise eluded detection by sensors and exited the game by reaching the target. Packet sizes are integral, so the probability is a discrete random variable.

b. Calculate Interdictor's Estimate of Flow on $(i, j) \in RS_{target}$

$\widehat{flow}_{(i,j),t}$: amount of material detected by a sensor on arc (i, j) during round t [flow/round]

$destroyed_{(i,j),t}$: total amount of material successfully destroyed on (i, j) by attack during round t [flow/round]

For $(i, j) \in RS_{"base"}$

$$\widehat{flow}_{(i,j),t} = \widehat{total.hits}_t + \sum_{(i,j)\in A} destroyed_{(i,j),t} - \sum_{(k,j)\in RS_{target} \cap k \neq i} \widehat{flow}_{(k,j),t} \quad (1.24)$$

Discussion

In (1.24), the interdictor forms an estimate of flow on each arc within the *RS* of the target. This estimate is based on the estimate of materiel exiting through the target in round t, $\widehat{total.hits}_t$, the amount destroyed in the arc (i, j) in round t, $\sum_{(i,j)\in A} destroyed_{(i,j),t}$,

and the total detected materiel on the arc (i, j) in round t, $\sum_{(k,j)\in RS_{target} \cap k \neq i} \widehat{flow}_{(k,j),t}$.

Uncertain of the actual path of undetected packets that have struck the target, the interdictor conservatively estimates that these packets travelled along every arc within the RS_{target}. This allows the interdictor to place sensors more reactively in the next round and mirrors realistic tactical decisions. Figure 31 provides an illustrative example of ***step 5.b***.

Figure 31. Interdictor's Estimate of Flow Reaching the Target.

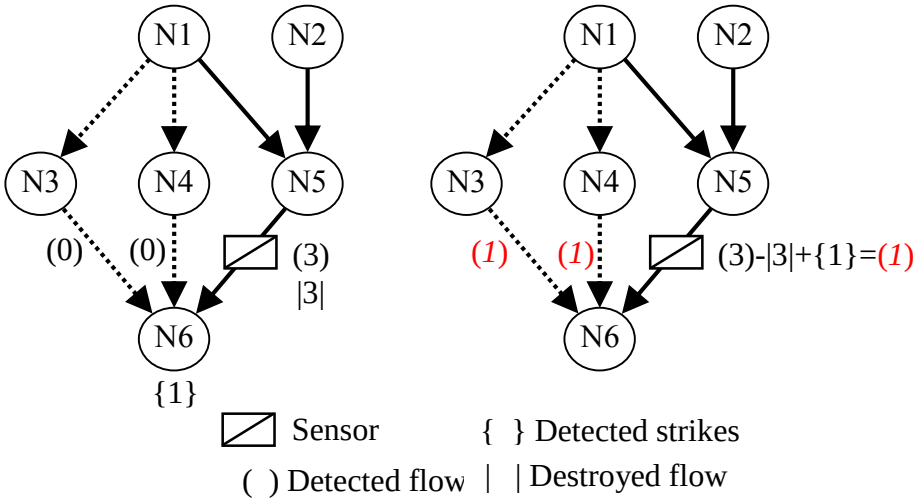

The observed flow on arc (*N5*, *N6*) is 3. The amount of destroyed flow on this arc is also 3. All other flows on the arcs in the *RS* of *N6*, the target, are 0. Each flow estimate is then updated by adding the amount of material striking the target and subtracting the amount of material destroyed on that arc. Values that change are highlighted in *red italics*.

c. Update Interdictor's Estimate of Arc Capacity and Cost Based on Detected Flow

γ: scaling parameter (0.0 – 1.0) that attenuates the interdictor's capacity estimate from the previous round, $\hat{u}_{(i,j),t-1}$.

For $(i, j) \in A$

$$\hat{u}_{(i,j),t} = \max\left(\left\lfloor \gamma \cdot \hat{u}_{(i,j),t-1} \right\rfloor, detections_{(i,j),t}, 1\right) \quad (1.25)$$

$$\hat{c}_{(i,j),t} = \min\left[\max\left(\hat{c}_{(i,j),t-1} - detections_{(i,j),t} + 1, 1\right), \hat{c}_{(i,j),t=0}\right] \quad (1.26)$$

If $\hat{u}_{(i,j),t} < \widehat{flow}_{(i,j),t}$ **Then** (1.27)

$$\hat{u}_{(i,j),t} = \widehat{flow}_{(i,j),t}$$

Discussion

In ***Step 6.c***., the interdictor *learns*, refining his estimate of arc capacities by utilizing the maximum of the amount of observed flow in round *t*, $detections_{(i,j),t}$, a scaled estimate of capacity from the previous round, $\gamma \cdot \hat{u}_{(i,j),t-1}$, and 1 (1.25). We apply the mathematical floor function to ensure this value remains an integer. The actual arc capacity must be at least the observed flow in round *t*. However, the interdictor reasons that detections might have been low or the smuggler could have utilized an alternative path. The interdictor retains some memory by considering the previous round's capacity, scaled by a factor, γ, between 0.0 and 1.0. Lastly, the interdictor always assumes each arc under consideration has capacity of at least 1. Even so, if the capacity estimate on arcs within RS_{target} is less than the flow estimated in ***Step 6.b*** (1.24), the interdictor conservatively increases the capacity estimate to the flow estimate (1.25).

By similar feedback, the interdictor updates arc cost estimates (1.26). The interdictor subtracts the number of detections in round *t*, $detections_{(i,j),t}$, from his cost estimate in the previous round, $\hat{c}_{(i,j),t-1}$. As with similar estimates, the interdictor adds 1 to the estimated arc cost so that decreasing detections on a particular arc cause the interdictor to look elsewhere for smuggled flow. Lastly, the interdictor places a lower bound on his current estimate of 1 and an upper bound equal to his initial round 0 estimate. This guarantees that the estimate of cost will become neither 0 nor excessively

large if the interdictor plays multiple rounds with a great many or no detections, respectively.

d. Update Smuggler's Arc Capacity and Cost Based on Game Round Results

attenuate: rate of cost attenuation [$\{attenuate \in \square \mid attenuate \in (0,1]\}$]

For $(i, j) \in A$

$$c_{(i,j),t} = \left\lfloor attenuate \times \left(c_{(i,j),t-1} + loss.multiple \times destroyed_{(i,j),t} \right) + 0.5 \right\rfloor \quad (1.28)$$

Discussion

Equation (1.28) instantiates a feedback loop that facilitates smuggler learning and resource loss. The smuggler adds the amount of material destroyed on an arc to the cost under the assumption that the loss of a packet increases the expense or risk of moving additional packets on the same arc. A *loss.multiple* magnifies this degradation to facilitate the exploration of various scenarios. The remaining elements of equation (1.28) instantiate the transformation function. We described them in the next section.

7. (Environment) Transformation Function

The transformation function first reduces the smuggler's arc cost by a constant factor, *attenuate*, and then rounds down to the nearest integer (1.28). For example, setting attenuate to 0.9 would reduce the cost on each arc $\in (i, j)$ by 10 percent before further rounding. The reduction in cost provides for repair of the network capability or "calming of risk" after the shock of a seizure by the interdictor. The reduction in cost also provides a mechanism to increase smuggler efficiency through time in accordance with traditional learning theory (e.g., Cesa-Bianchi et al. 2006).

E. DISCUSSION OF I-I/S MODEL FORMULATION

The preceding **I-I/S** mathematical model formulation displays both the family of realistic features and suite of complimentary stochastic and optimization models we bring to bear in order to study the interdictor-smuggler problem. Our modelling assumptions—asymmetric, incomplete information between two players receiving limited feedback over multiple rounds of play—would make finding optimal solutions computationally

prohibitive if not entirely intractable for even many of the most modern models reviewed within Chapter 2 of this thesis. However, by rigorously blending a number of well-established analytic tools within our simulation algorithm, we provide a tractable heuristic to explore a rich, complex problem under limited assumptions.

In the next chapter, we continue the exploration by applying the **I-I/S** simulation algorithm to realistically configured test cases.

IV. RESULTS AND ANALYSIS

In this chapter, we present the results from two *interdictor-smuggler* problem instances. The first instance portrays a smuggling network with a large number of possible parallel routes, the second a network with fewer alternative paths but significant depth. Our goal is to assess the performance of various interdictor resource allocation policies and illustrate the types of insights that our model can provide.

A. COMPUTATIONAL CASES

1. Implementation Details

We code the **I-I/S** algorithm in a combination of Python 2.7 (Van Rossum 2007) and GAMS 24.4.1 (GAMS Development Corporation 2013). Within Python 2.7, we utilize the NetworkX package (Hagberg et al. 2008). We use GAMS with CPLEX 12.6.1.0 to solve the SENSOR_PLACEMENT Dual Integer-Linear Program (ILP) and MOVE_PACKETS ILP formulated in Chapter 3. Sanchez's NOLHdesigns spreadsheet was used to construct the two Orthogonal Latin Hypercube designs of experiments (Sanchez 2015). We perform all experiments on a Windows PC with 2.6 GHz CPU and 8 GB RAM. Lastly, we use a combination of R 3.3.0 (R Core Team 2016) and JMP Pro 12.0.1 (JMP Pro 2015) to construct the statistical models herein.

2. Case 1

Case 1 involves a smuggling network over 27 arcs and 16 nodes (Figure 32). The designed network provides 25 paths from the smuggler's source of materiel, node *HB* (*Hostile Base*), to the intended target, node *N13*. We place arc (*N13, N15*) within the network to deliberately explore instances of miscalculation by the interdictor. As it is not possible for the smuggler to ever move materiel across arc (*N13, N15*), the interdictor would be wrong to place sensors there. Table 4 displays the initial network data. The smuggler's estimates of arc capacities begin as set uniformly to 10 units in order to reduce the confounding of dynamic cost and capacity effects under analysis. At the initiation of the game, the interdictor considers each route as equally likely. Therefore,

similar to the smuggler's estimates of arc capacity, the interdictor's estimates of arc costs and capacities are originally uniform. We then design the network topology and initial smuggler cost estimates to create a realistic range of total route costs and facilitate the introduction of new routes over successive game rounds.

Figure 32. Case 1 Designed Network.

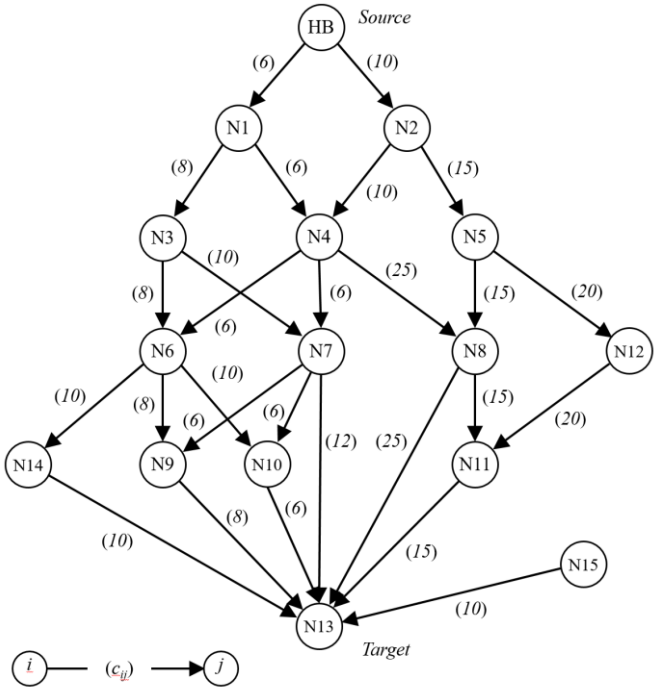

There are 25 *s-t* paths in the above network. Initial smuggler estimates of arc cost are indicated in parenthesis. Smuggler arc capacity estimates are uniformly 10.

We consider a time horizon of 20 game rounds for all scenarios in Case 1. Prototype simulation runs within the Case 1 problem instances demonstrated that within 20 game rounds we observe several cycles of adaptive play without incurring excessive computational expense. Dynamic topology motivates both the smuggler and interdictor to adapt. In Case 1, all arcs with cost above 8 units are initially invisible to the smuggler. The selected value, 8, allows the transformation function to reveal 15 new arcs to the smuggler within the 20 game rounds. These new arcs progressively introduce 11 of the 25 total paths throughout each 20-round game (Figure 33). Given our selected network

data, we find that both the smuggler and interdictor will adjust to a new path within approximately three to four rounds of play.

Table 4. Initial Interdictor and Smuggler Estimates.

Arc		Interdictor Estimates		Smuggler Estimates	
Tail	**Head**	**Cost**	**Capacity**	**Cost**	**Capacity**
HB	N1	10	1	6	10
HB	N2	10	1	10	10
N1	N3	10	1	8	10
N1	N4	10	1	6	10
N2	N4	10	1	10	10
N2	N5	10	1	15	10
N3	N6	10	1	8	10
N3	N7	10	1	10	10
N4	N6	10	1	6	10
N4	N7	10	1	6	10
N4	N8	10	1	25	10
N5	N8	10	1	15	10
N5	N12	10	1	20	10
N6	N9	10	1	8	10
N6	N10	10	1	10	10
N7	N9	10	1	6	10
N7	N10	10	1	6	10
N8	N11	10	1	15	10
N8	N13	10	1	25	10
N12	N11	10	1	20	10
N9	N13	10	1	8	10
N10	N13	10	1	6	10
N11	N13	10	1	15	10
N14	N13	10	1	10	10
N7	N13	10	1	12	10
N6	N14	10	1	10	10
N15	N13	10	1	10	10

The Case 1 initial network data. In order to reduce the confounding of dynamic cost and capacity effects under analysis, the smuggler's estimates of arc capacities begin as set uniformly to 10 units. The interdictor considers each route as equally likely at the initiation of the game. Therefore, similar to the smuggler's estimates of arc capacity, the interdictor's estimates of arc costs and capacities are originally uniform. We then design the network topology and initial smuggler cost estimates to create a realistic range of total route costs and facilitate the introduction of new routes over successive game rounds.

Figure 33. Initial Arcs Visible to the Smuggler and Interdictor

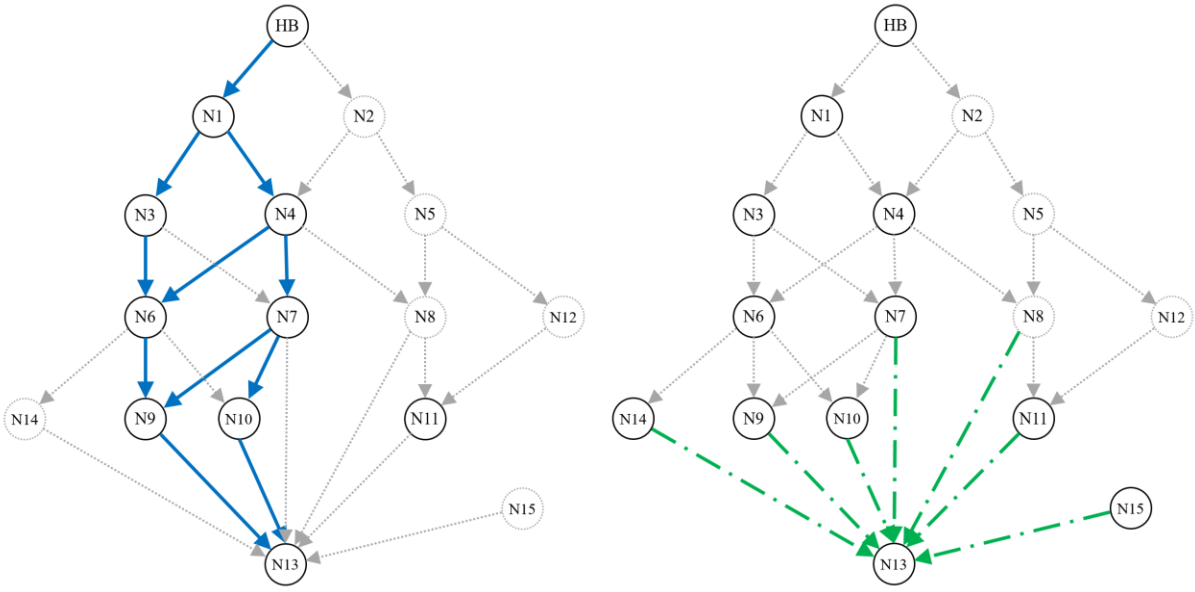

Left: Arcs visible to the smuggler in round 1 (heavy blue). The remaining arcs are revealed by the transformation functions action as a timer. *Right*: Arcs visible to the interdictor in round 1 (dashed green). Only arcs adjacent to the interdictor's placed sensors are visible. The interdictor remembers all arcs seen in previous rounds. In Case 1, we observe the two subgraphs seldom match.

The smuggler's budget, 350 cost/round, is just adequate to allow the smuggler to transport any materiel introduced at node *HB* in a single round through the network within the two following rounds. Requiring a minimum of two rounds to traverse the network allows us to investigate interdiction decisions that must consider both time and space. Larger budgets that provide sufficient resources for the smuggler to move flow across the entire network within only one round create a problem only in space, not time, for the interdictor. Without interdiction in the Case 1 network, the maximum flow per round is limited by budget, not arc capacity. The smuggler could transport 20 units of flow through the entire network within one round, given a sufficiently large movement budget. Therefore, representing a situation where the smuggler's budget is just sufficient to accomplish his immediate aim of moving a packet from source to target in two rounds also allows a more transparent quantification of the impact of the interdictor's resource allocation decisions.

In Case 1, the master scheduling table introduces 240 total units of materiel at node *HB* across the 20 rounds. The materiel is divided into packets either uniformly-sized at size 1, 2 or 3 or randomly-sized between 1 and 3 units. A Mersenne Twister algorithm generated the sequence of pseudo-random packet sizes (Matsumoto 1998). We partition the resultant sequence of packet sizes into 20 groups so that each partition includes approximately 12 units of flow (Table 5). There are approximately 12 units of flow in each round whether the packets are all 1, 2, 3, or randomly-sized. The smuggler's supply is thus consistent in each round.

Table 5. Example of Master Packet Schedule.

Round	Packet	Size	Round	Packet	Size
1	P1	3	2	P10	1
1	P2	2	2	P11	3
1	P3	2	2	P12	3
1	P4	1	3	P13	3
1	P5	1	3	P14	1
1	P6	2	3	P15	2
2	P7	2	3	P16	3
2	P8	2	3	P17	1
2	P9	2	3	P18	2

In some scenarios, packet sizes are randomly drawn between one and three using the Mersenne Twister algorithm. The sequence of random packet sizes is partitioned to introduce approximately 12 units of flow to each round. However, the table continues to round 20; only rounds 1–3 are displayed in the above example.

a. Design of Experiments

We combine an Orthogonal Latin Hypercube (OLH), crossed design, and star points to consider 71 equally likely scenarios in Case 1. These scenarios program two decision variables and five noise variables through a realistic range of values (Table 6).

Table 6. Case 1 Range of Factors in Experimental Design.

Decision Variables	**Definition from Chapter 3**	**Range of values**
$covert.budget_t$	maximum number of covert sensors the interdictor can place in round *t* [cardinality/round]	[0, Sensor Budget]
$overt.budget_t$	maximum number of overt sensors the interdictor can place in round *t* [cardinality/round]	[0, Sensor Budget]
Noise Factors	**Definition from Chapter 3**	**Range of values**
$q_{(i,j),t}$	smuggler penalty for traversing an arc with sensor [cost/flow-round]	[10, 50]
$size_p$	size of packet *p* [flow]	[1, 2, 3, random integer [1,3]]
packet.signature	level of stealth for packets [non-negative integer]	[1, 2, 3]
γ	scaling parameter (0.0 – 1.0) that attenuates the interdictor's capacity estimate from the previous round, $\hat{u}_{(i,j),t-1}$	[0.6, 1.0]
attenuate	rate of cost attenuation $[\{attenuate \in \square \mid attenuate \in (0,1]\}]$	[0.88, 0.95]

The decision variables represent choices the interdictor makes to allocate his sensor budget between covert and overt sensors, grossly described by three general policies: *pure overt*, *pure covert*, and *hybrid* policy. Under pure overt and pure covert policies, the interdictor devotes the entire sensor budget to overt or covert sensors, respectively. The hybrid policy includes each possible combination of overt and covert sensor allocation within the overall sensor budget. For example, hybrid policies within sensor budget 3 include both (a) 1 overt and 2 covert sensors, and (b) 2 overt and 1 covert sensors. The range of values assumed by the noise variables in Table 6 acts as a surrogate for different levels of smuggler and interdictor technological sophistication and tactical acumen. We use a two-stage statistical sampling procedure as developed by Dudewicz

and Dalal (1975) to determine the number of required replications. Under the procedure, we repeat each scenario a minimum of 10 times to establish at least a 90 percent confidence level that the error in an estimate of the expected value of smuggled flow is less than 10 percent.

b. Results

(1) Interdictor Policy Performance: Aggregate Metrics

We consider any smuggled flow reaching the target as *loss*. Figure 34 shows the total smuggled flow (loss) versus the interdictor's sensor budget. The plot clearly shows a face-valid trend: with additional sensors, the interdictor can cause greater disruption to smuggler routes, reducing smuggler flow. We seek to understand this relationship in more detail. (For further information on robust design and loss functions see Kleijnen et al. [2005], Sanchez [2000].)

Figure 34. Total Smuggled Flow versus Interdictor Sensor Budget.

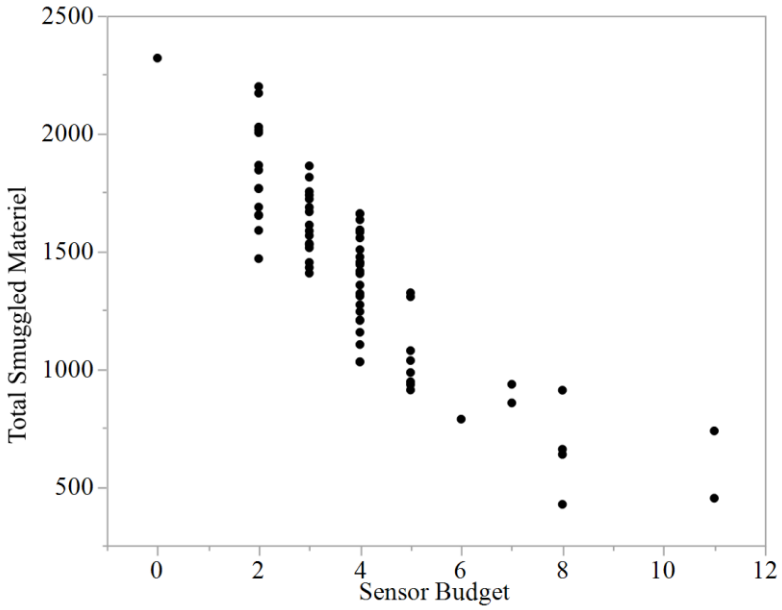

The total smuggled flow per game decreases as the interdictor's sensor budget increases. The lack of data at sensor budgets 6, 9 and 10 is an artifact of the OLH design. Because the results from each policy show unequal variance, we use a loss function to calculate the total smuggled flow and then compare policy performance.

Within the range of scenarios described in Case 1, we further focus our investigation on instances where the interdictor has a sensor budget between two and five. Given a budget of six or more sensors, the interdictor can proceed with a naïve strategy, placing sensors along a majority of the always-visible *s-t* cut across the reverse star of node *N13* (Figure 35). Such a strategy requires neither estimation nor adaptation from the interdictor to execute successfully. Within our model, the smuggler's private and incomplete information provides a subtle nuance for the interdictor's strategic decisions. Given that the smuggler plays on a subset of the available arcs that changes in time, the minimum *s-t* cut of the subgraph created by these arcs also changes in time. Restricting the interdictor to less than six sensors thus allows us to investigate problem instances where there is no obvious sensor allocation strategy that is robust to an adapting and malicious smuggler.

Figure 35. A Naïve *s-t* Cut in Case 1.

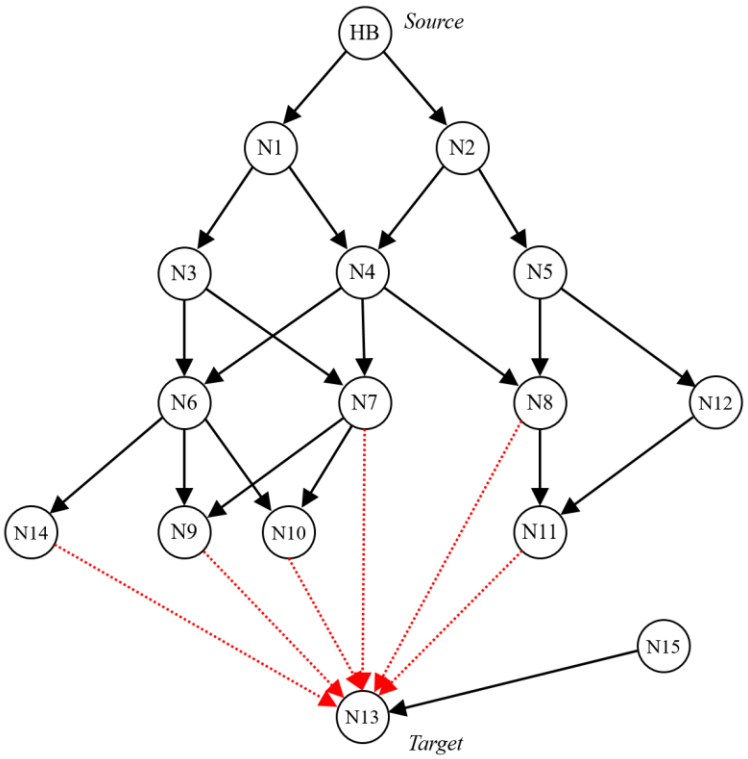

Six sensors are sufficient to make an *s-t* cut (dashed red) with no exploration from the interdictor. Because they are part of the target's reverse star, these arcs are always visible to the interdictor. We limit our further investigations to interdictor sensor budgets below six to examine the value of less obvious interdictor strategies.

We supplement our original OLH design with 53 new design points. Using an augmented crossed design, the experiments consider the same range of noise factors for sensor budgets restricted between two and five, inclusive (Table 7). The additional 53 design points create a much higher degree of space filling and allow verification of the statistical model predictions we generate using the original OLH design. Because the variance of each studied instance is unequal and variance is generally undesirable in tactical situations, we use a loss function to compare policy effectiveness. Figure 36 displays the loss by policy. Figure 37 displays the percentage of flow degradation by policy for each sensor budget.

Table 7. Case 1 Range of Factors in Crossed Design with Star Points.

Decision Variables	**Range of values**
$covert.budget_t$	[2, 5] | $covert.budget_t + overt.budget_t$ = [2,5]
$overt.budget_t$	[2, 5] | $covert.budget_t + overt.budget_t$ = [2,5]
Noise Factors	**Range of values**
$q_{(i,j),t}$	[20, 40]
$size_p$	random integer [1,3]
packet.signature	[1, 3]

The interdictor desires to minimize the smuggled flow. Given two or three sensors, allocating them all as overt sensors, a *pure overt* policy, minimizes flow. These policies are 8 percent and 3 percent better than the hybrid policies with one covert sensor, the next best performing policies (Table 8). However, as the sensor budget increases to four and five, the hybrid policy becomes more effective and eventually superior to the pure overt policy (Figure 36, Table 8).

We apply 10 additional design points across each sensor budget. The added design points allow finer resolution of the near equal loss resultant from several policies. We note that at sensor budget four, the hybrid policy and pure overt policy with one

covert sensor produce almost the same degradation of flow, 49 and 50 percent, respectively (Table 8). However, when we consider total materiel destroyed, there is a significant difference in policy performance (Figure 37). At sensor budget five, the amount of seized materiel is also significantly different between the two policies with maximum degradation of smuggled flow, hybrid (3 covert) and hybrid (4 covert).

Figure 36. Contour Profile of Loss versus Number of Overt and Number of Covert Sensors.

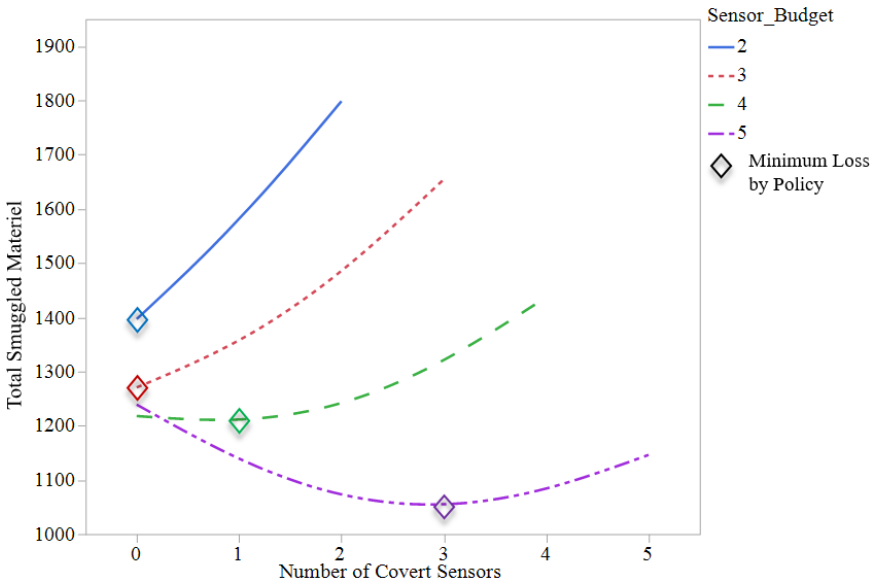

The graphic shows loss contours projected from our statistical model (Appendix A). The contours indicate *pure overt* sensor policies produce the minimum loss for sensor budgets equal to two or three. However, the *hybrid* policy is a superior allocation of either four or five total sensors.

Table 8. Policy Performance: Percentage of Degradation of Smuggler Flow.

	Number of Covert Sensors					
Sensor Budget	0	1	2	3	4	5
2	***42%***	34%	25%	NA	NA	NA
3	***47%***	44%	38%	31%	NA	NA
4	**49%**	***50%***	48%	45%	40%	NA
5	48%	53%	**55%**	***56%***	**55%**	52%

Policy: *Pure Overt* | *Hybrid* | *Pure Covert*

At sensor budgets equal to two and three, pure overt policies maximally degrade the smuggler's flow. However, hybrid policies become more effective at with a budget of four sensors, and clearly superior with five total sensors available.

Figure 37. Policy Performance by Sensor Budget.

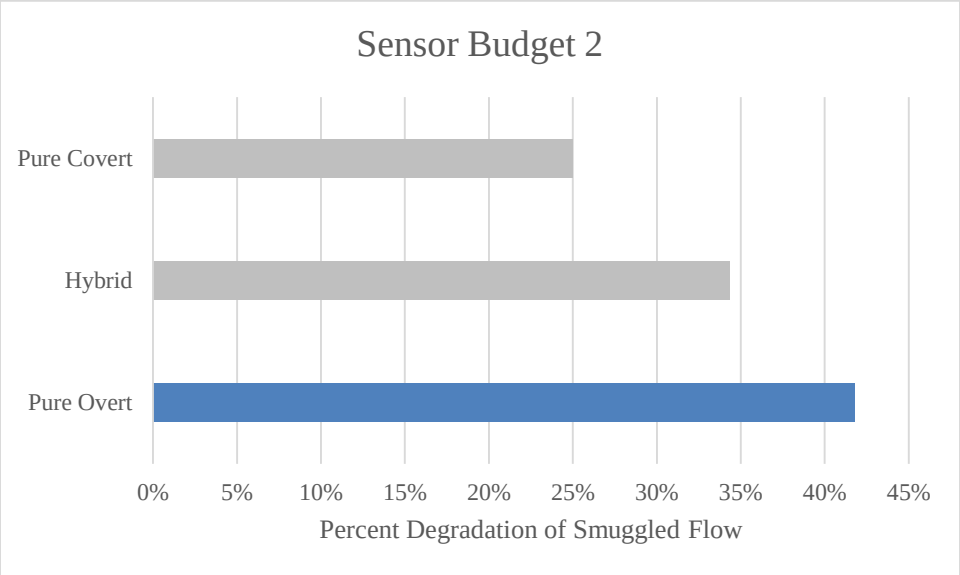

(a) With a sensor budget equal to two, the pure overt strategy maximally degrades the smuggled flow. It offers 8% more degradation than the hybrid policy.

Figure 37. Policy Performance by Sensor Budget (Cont.).

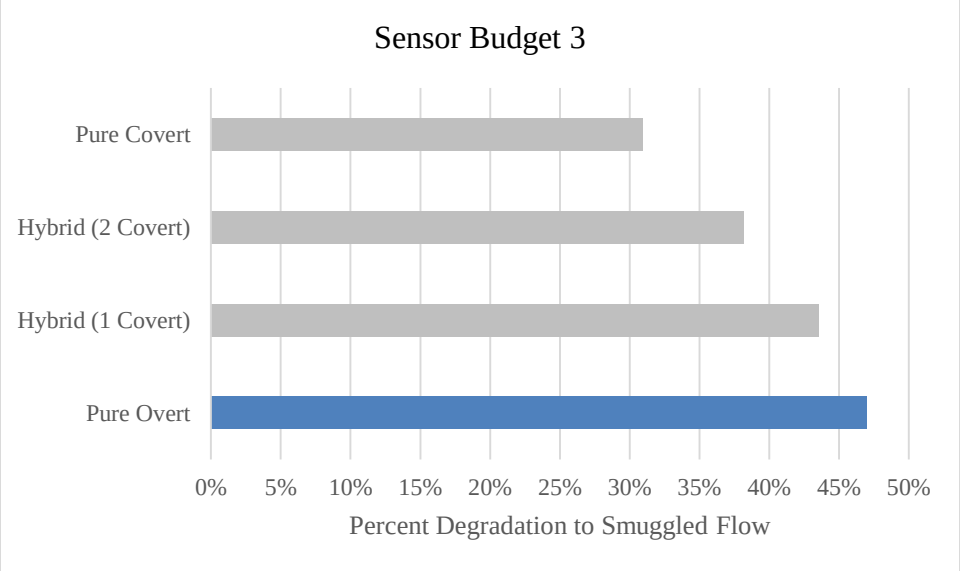

(b) The pure overt policy still offers maximum degradation of smuggled flow with three sensors available. The degradation is 3% greater than the hybrid (1 covert) policy.

Figure 37. Policy Performance by Sensor Budget (Cont.).

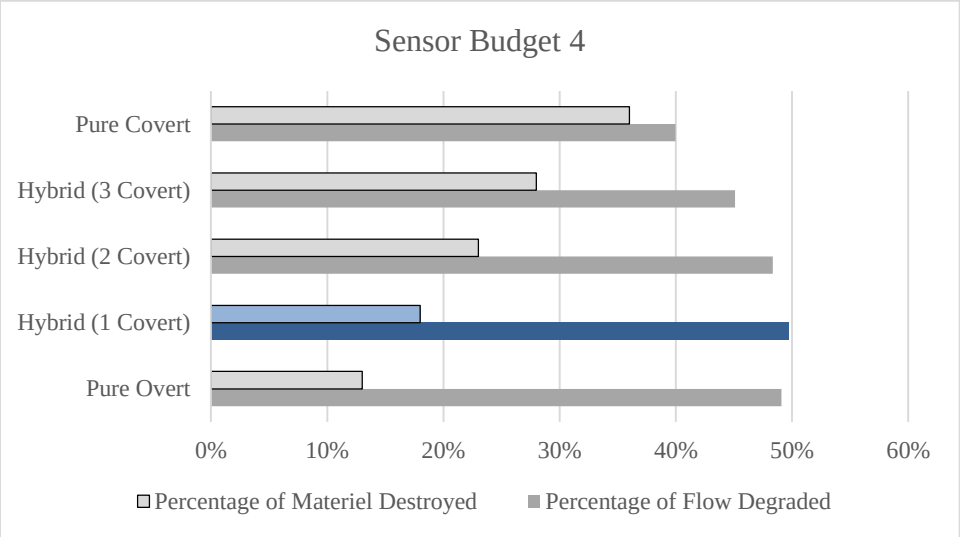

(c) While the pure overt and hybrid (1 covert) policy showed the almost the same degradation of smuggled flow, 49%, there is significant difference in the percentage of total materiel destroyed. The hybrid (1 covert) policy destroys over 6% more flow (darker in blue).

Figure 37. Policy Performance by Sensor Budget (Cont.).

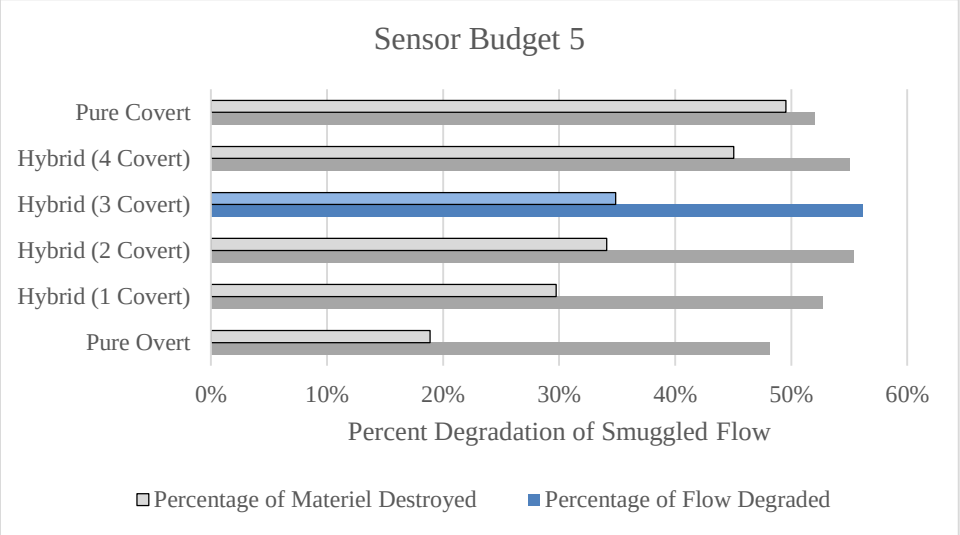

(d) The hybrid (3 covert) policy maximally degrades the smuggler's flow, given five available sensors. Under the hybrid policy, the degradation is 2% greater than the next highest option, a hybrid (4 covert) policy. Even so, the amount of materiel destroyed increases by 10% when the interdictor follows a hybrid (4 covert) policy.

Within the detail provided by the expanded experimental design, we also note a significant improvement in flow degradation as sensor sensitivity improves. Figure 38 shows the effect on the percent of smuggled flow degraded versus the number of covert sensors as sensor probability of detection increases by approximately 8 percent. An 8 percent change encompasses the full range we examine in our experimental design. The improved degradation of flow was consistent across all Case 1 policies. As the number of covert sensors increase, the improvement ranged from approximately 7 percent to 24 percent degradation of flow. Even so, the improvement did not significantly change until the interdictor employed more than two covert sensors.

Figure 38. Effect of Sensor Sensitivity on the Percent Degradation of Smuggled Materiel.

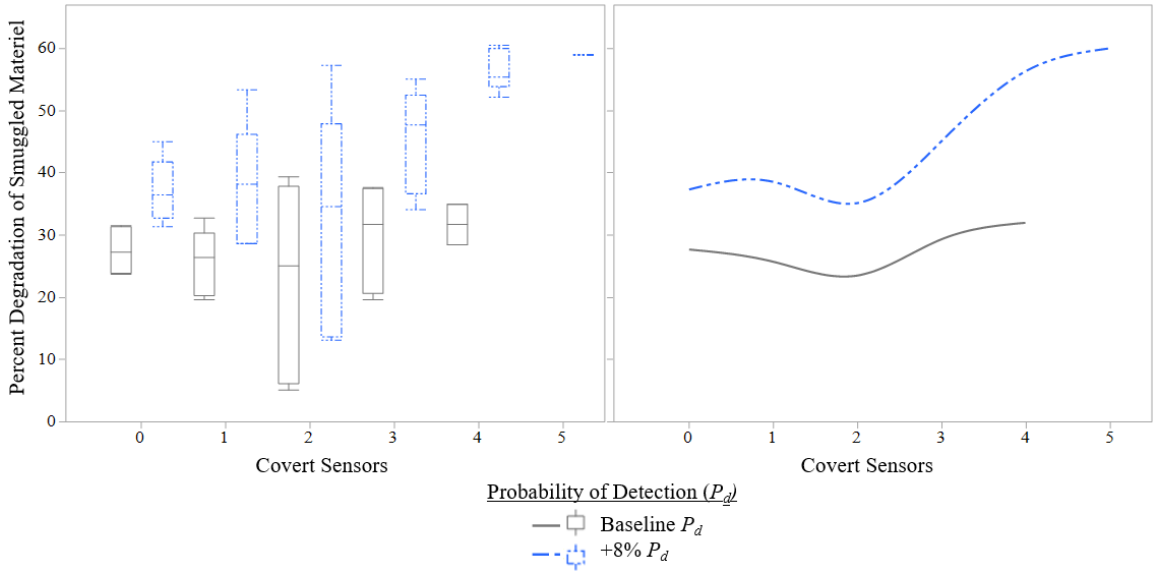

Improving each sensor's probability of detection by approximately 8% yields consistent improvement in all Case 1 scenarios. *Left*: Boxplots illustrate the general distribution of the percent of degradation of smuggled materiel versus the number of covert sensors by probability of detection scenario (Baseline P_d or +8% P_d). *Right*: To ease interpretability and comparison, we plot two statistical models (piecewise splines with $\lambda = 0.05$), one fit to the Baseline P_d and one fit to the +8% P_d. The 8% increase to probability of detection causes approximately 7% greater degradation of flow with low numbers of covert sensors. As the number of covert sensors increases, so does the effect of increased P_d. With four covert sensors, there is a 24% increase to smuggler flow degradation when the probability of detection is raised approximately 8%.

(2) Interdictor Policy Performance: Time Dynamics

Having considered the aggregate performance of each policy, we turn our attention to the detailed time dynamics of game play under each policy. We use the 710 time series generated from our DOE to construct a *meta-game* for each policy, representing the 0.9 quantile of per-round smuggled flow of all games in the policy. That is, only 10 percent of the games had higher per-round smuggled flow than the meta-game. Figure 40 depicts the family of meta-games within each sensor budget. In that figure as a reference, we highlight the policy generating the greatest overall degradation to the total smuggled flow, as analyzed in Case 1, Section A.2.b.(1). It is immediately apparent that the games are highly dynamic. Round-to-round flow oscillates significantly. Further analysis of smuggler packet movement and interdictor sensor placement shows the oscillation is a result of the action-counteraction cycle between the interdictor and smuggler aggravated by the revelation of new arcs through the countdown timer. We also note that certain policies tend to reduce more consistently the maximum per-round smuggler flow within a game. We focus on round 8, when by the action of the countdown timer, the simulation makes an entirely new smuggling path visible (Figure 39).

Figure 39. Newly Visible Smuggling Path in Round 8.

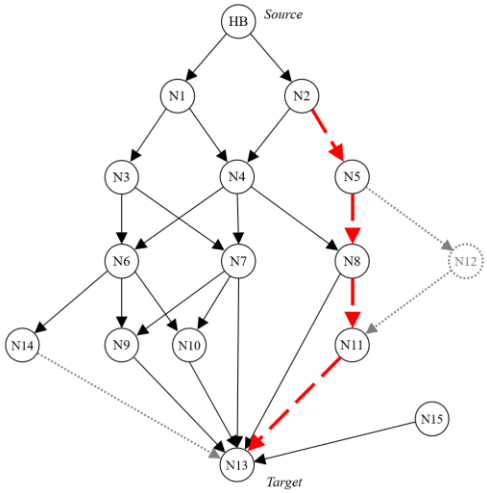

The countdown timer reveals the *N2—N5—N8—N11—N13* path in round 8 (dashed red), almost an entire *s-t* path. Arcs (*N5, N12*), (*N12, N11*), and (*N14, N13*) are not yet visible (dotted).

Figure 40. *Meta-Games*, the 0.9 Quantile of Smuggled Flow by Game Round.

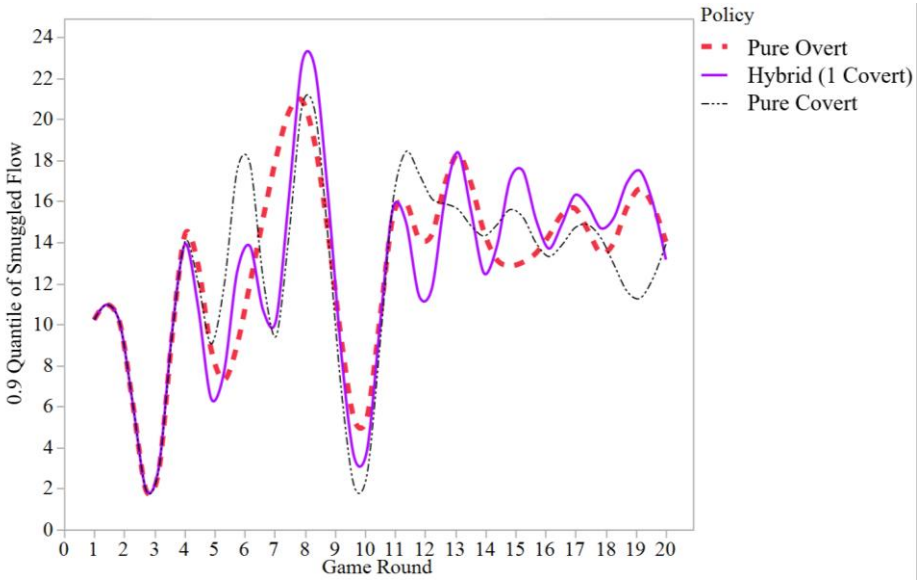

(a) **Sensor Budget 2**. The smuggled flow dramatically oscillates round-to-round by almost an order of magnitude. No single policy shows significantly different performance than the pure overt policy (dashed red). Flow peaks at 20 to 24 units in round 8 when the algorithm introduces a new *s*-*t* path, but the interdictor rapidly adjusts by round 10 dropping flow to 4 units. Heavy red dashes indicate the policy maximally degrading the total smuggled flow per game in Case 1.

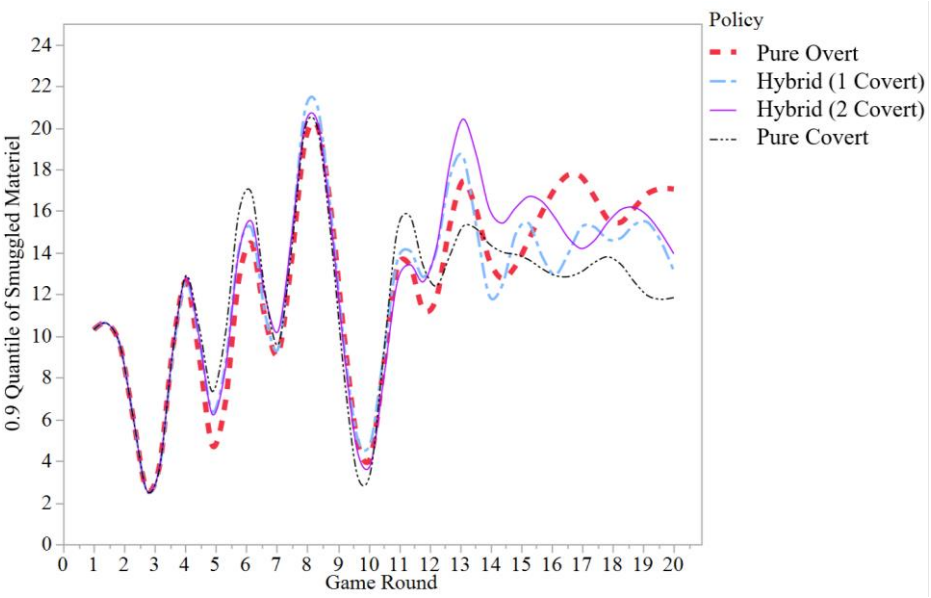

(b) **Sensor Budget 3**. As with a budget of two sensors, given three sensors, we observe no significant difference between each policy's round-to-round play. The peak flow in round 8 is slightly attenuated, reducing from 20–24 to 20–22. Heavy red dashes indicate the policy maximally degrading the total smuggled flow per game in Case 1.

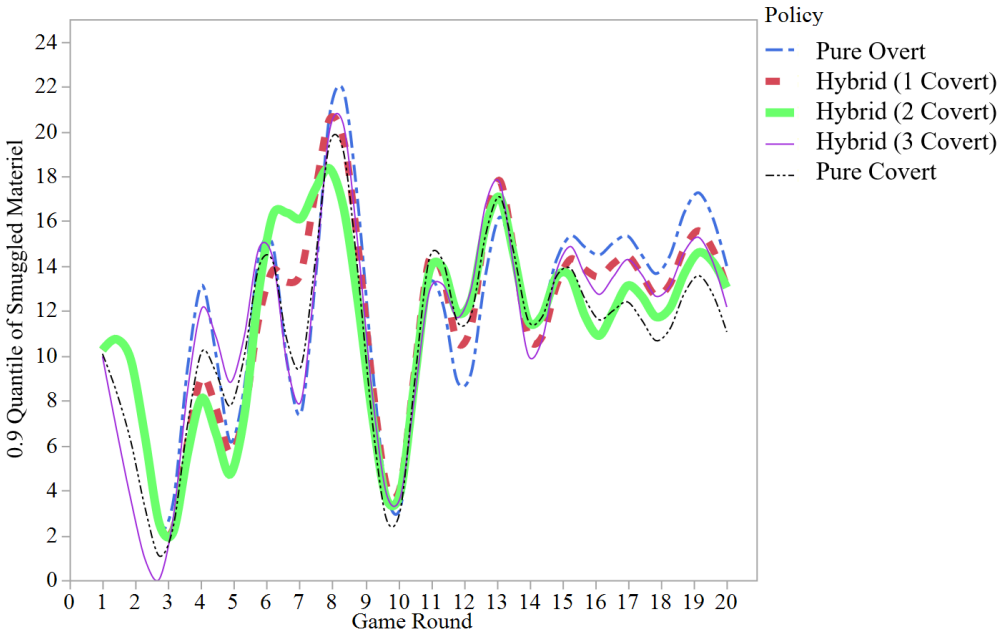

(c) **Sensor Budget 4**. We begin to observe significant differences in policy round-to-round performance. Of note, the hybrid (2 covet) policy, depicted in green, shows a 10–20% reduction in the round 8 peak flow. Red dashes indicate the policy maximally degrading the total smuggled flow per game in Case 1.

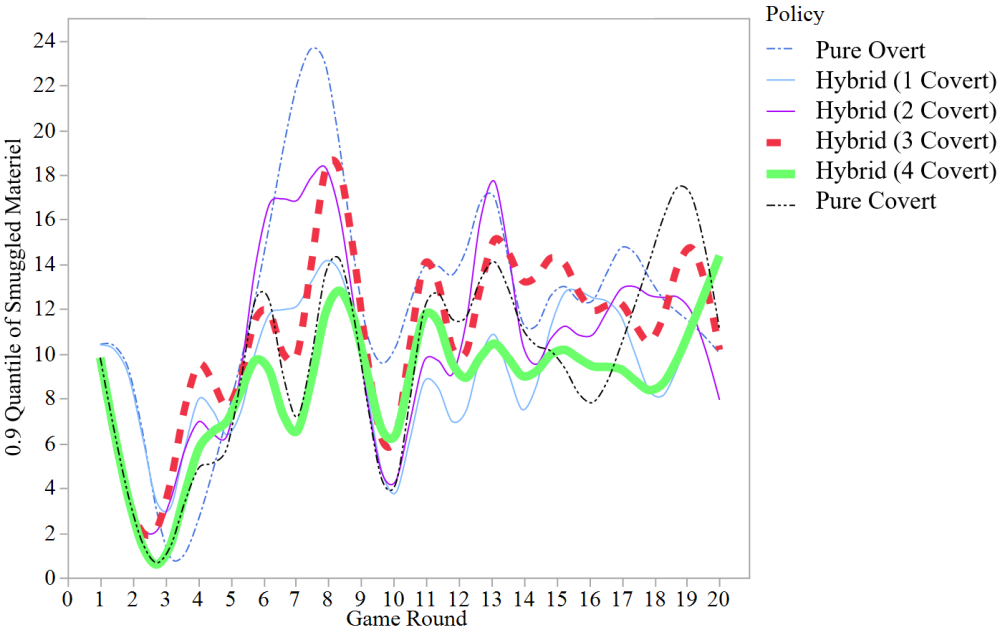

(d) **Sensor Budget 5**. There is now a substantial difference in round-to-round policy performance. The hybrid (4 covert) policy demonstrates a reduction in round 8 peak smuggled flow of over 50 percent. A 20–30% reduction in the peak flow of later rounds is also apparent. Red dashes indicate the policy maximally degrading total smuggled flow in per game Case 1.

As sensor budgets increase, we observe significant differences by policy in peak round-to-round smuggled flow. The 0.9 quantile of smuggled flow reduces by up to 50 percent in round 8 when the interdictor employs a hybrid policy over pure overt policy.

(3) Evaluation of Seizures as Proxy for Smuggled Flow

Figure 41 shows the total smuggled flow versus the total amount of materiel destroyed over all Case 1 scenarios. A naïve regression indicates a moderately strong relationship between the total amount of materiel seized and the total amount of smuggled materiel (R^2 = 0.605).

Figure 41. Linear Regression of Total Smuggled Materiel versus Total Materiel Detected.

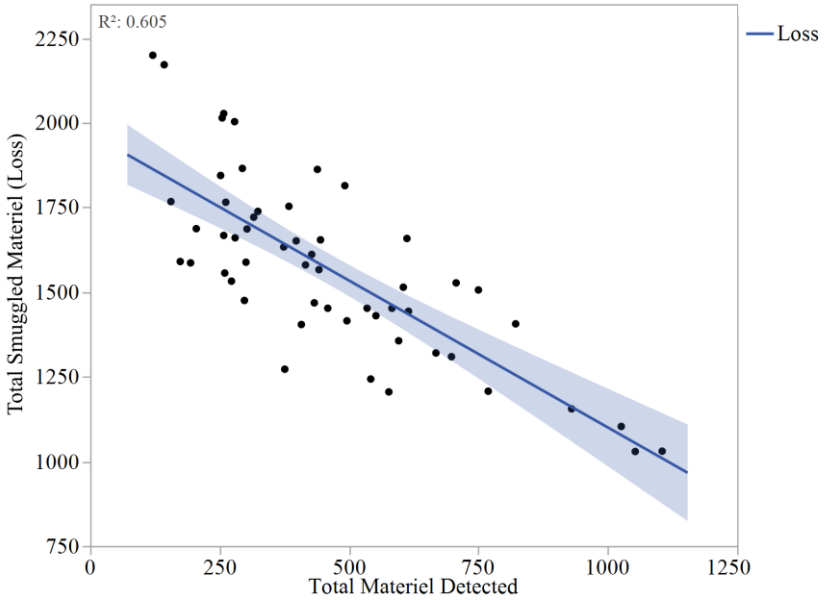

The best linear model explains only 60 percent of the observed variance in the total smuggled materiel. We observed no improved explanatory power under a variety of variable transformations.

However, within the naïve regression, we note an inconsistent level of variability indicating heteroscedasticity, bringing doubt on any inference from the regression results (Figure 42). Further exploration indicates a significant change in the strength of the relationship between seizures and flow when total sensor budget is considered (Figure 43). With a budget of three sensors, the amount of seized materiel explains only 23

percent of the variability in the amount of unseen smuggled materiel. Even so, models considering sensor budget still suffer from some inequality of variance at different simulation configurations representing the different interdictor policies.

Figure 42. Residuals for the Naïve Regression Model.

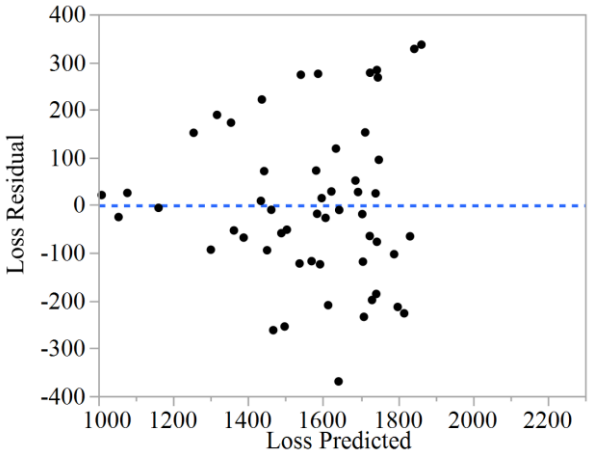

The plot of residuals indicates unequal variance among estimates. The heteroscedasticity makes any inference problematic.

Figure 43. Total Smuggled Materiel versus Total Materiel Detected.

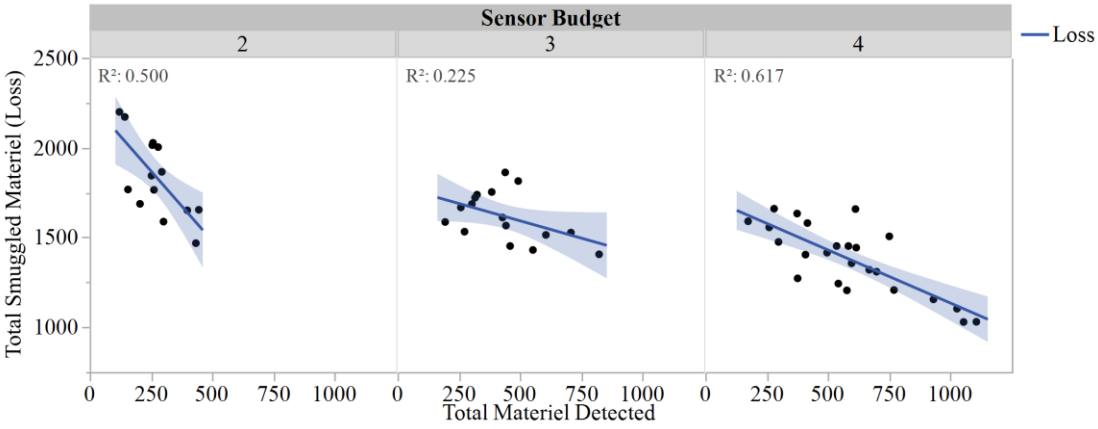

The relationship between total materiel detected and total materiel smuggled is inconsistent. *Left*: With a budget of two sensors, the amount of detected materiel explains only 50 percent of the variance in the total smuggled flow. *Center*: With a budget of three sensors, the relationship between detected flow and actual total flow is extremely weak. *Right*: The largest R^2 occurs with a budget of four sensors, yielding 62 percent of the variance explained. The shaded region represents a 95 percent confidence interval.

When we also consider individual policies, variance stabilizes, providing a much stronger case for accurate inference (Figure 44, Figure 45).

Figure 44. Model of Total Smuggled Materiel versus Total Materiel Detected by Policy with a Budget of Three Sensors.

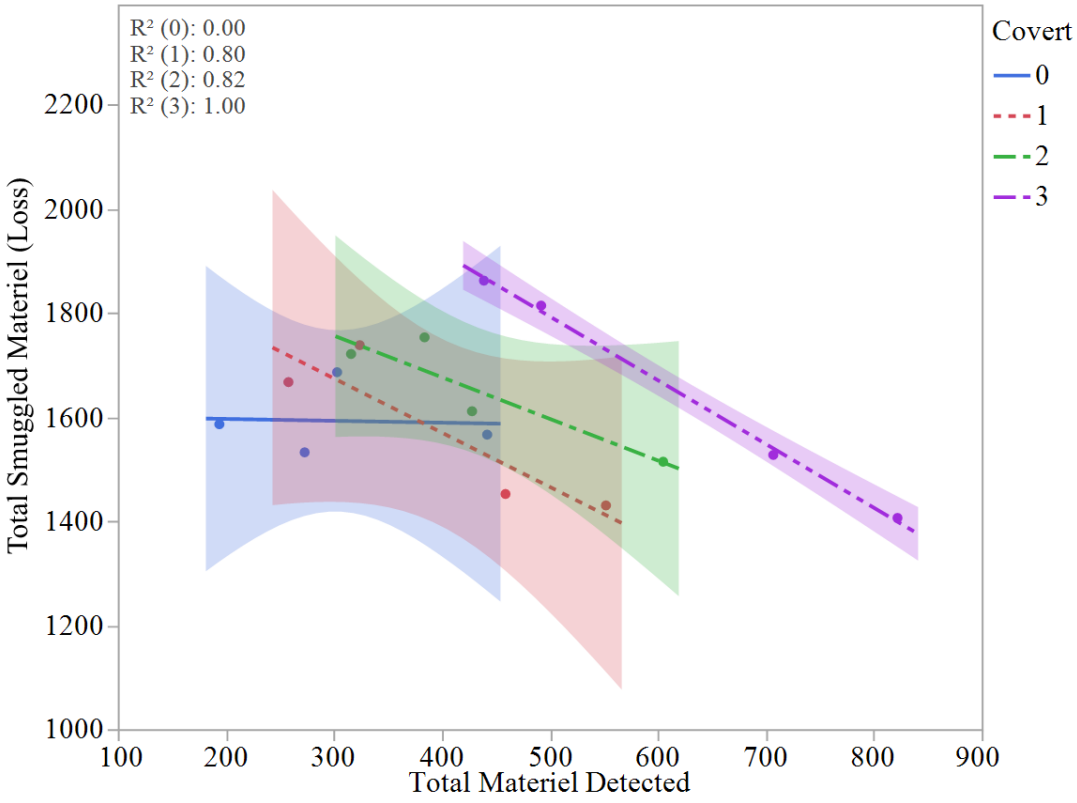

The relationship between the total materiel detected and the total smuggled materiel is highly inconsistent. For the pure overt policy (blue), the total materiel detected conveys no information on the total smuggled flow, R^2 = 0.00. The shaded regions represent a 95 percent confidence interval.

Figure 45. Model of Total Smuggled Materiel versus Total Materiel Detected by Policy with a Budget of Four Sensors.

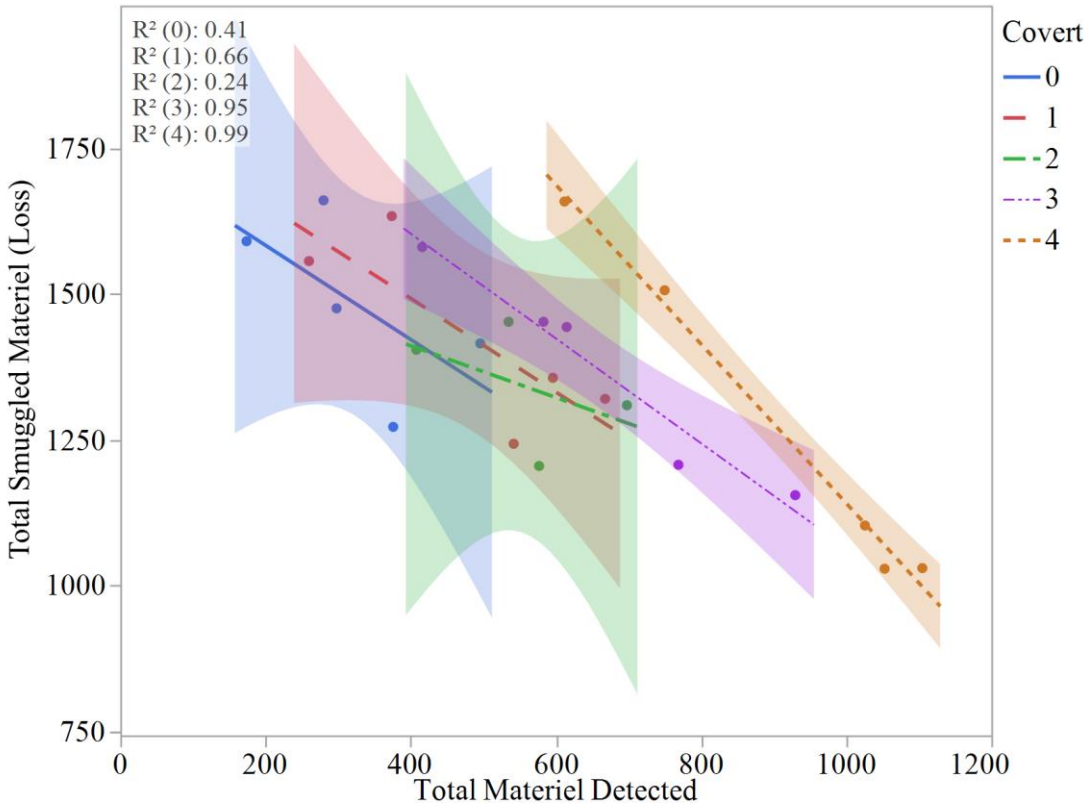

With a budget of four sensors, the relationship between the total materiel detected and the total smuggled materiel is again highly inconsistent. It varies from high explanatory power for pure covert policy (dashed orange) to very low explanatory power for hybrid (2 covert) (green). The shaded regions represent a 95 percent confidence interval.

As shown in Figure 44 and Figure 45, there is wide variation in the power of the amount of detected materiel to explain the amount of smuggled flow varies considerably. With pure covert policies, the relationship is extremely strong. However, for three policies featuring more overt sensors, the amount of detected materiel is weakly related to the total smuggled materiel. Generally, these same more overt policies are also the policies we found most effective at minimizing the total smuggled materiel.

(4) The Relationship between Discovered Network and Flow

In this section, we explore the relationship between the information gained by the interdictor through exploration of the smuggling network and the total amount of

smuggled flow. First, to quantify the information gained by exploration, we measure the percentage of arcs within the smuggling network discovered by the interdictor under each sensor allocation policy. Displayed in Figure 46, policies with higher percentages of covert sensors tend to discover more of the smuggling network but do not explain a large portion of the variance for the network discovered.

We continue assessing the value of information by now quantifying it by the total amount of materiel smuggled. Figure 47 shows the total smuggled materiel versus the percentage of all arcs discovered by the interdictor within each of the 47 Case 1 scenarios. Using a generalized linear model, the percentage of arcs discovered only explains 10 percent of the variance in the total amount of materiel smuggled. Further investigation by specific interdictor policy also reveals weak relationships between these two factors.

Figure 46. Percent Arcs Discovered versus Percent of Sensor Budget Dedicated to Covert Sensors.

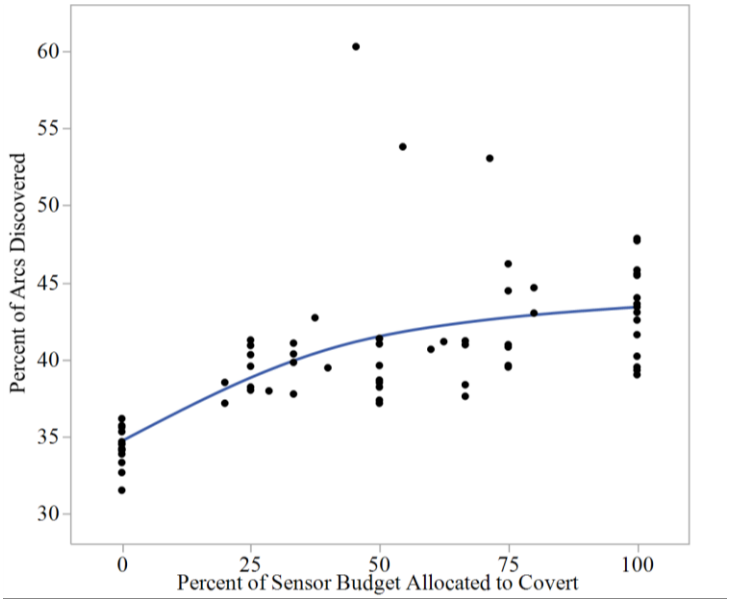

The above data results from measuring with a loss function the percent of total arcs discovered. The policies have heterogeneous variance. The loss function allows us to compare different policies despite the heterogeneity. Interdictor policies with higher percent of covert sensors generally discover more of the smuggler network. Even so, the relationship is weak, accounting for just over 40 percent of the variance in the nonparametric kernel density smoother using local weights (blue line). We draw attention to this face-valid trend to support observations focused on less-obvious emergent

relationships in the following sections. (piecewise splines with $\lambda = 0.05$),

Figure 47. Percentage of Arcs Discovered versus Total Loss.

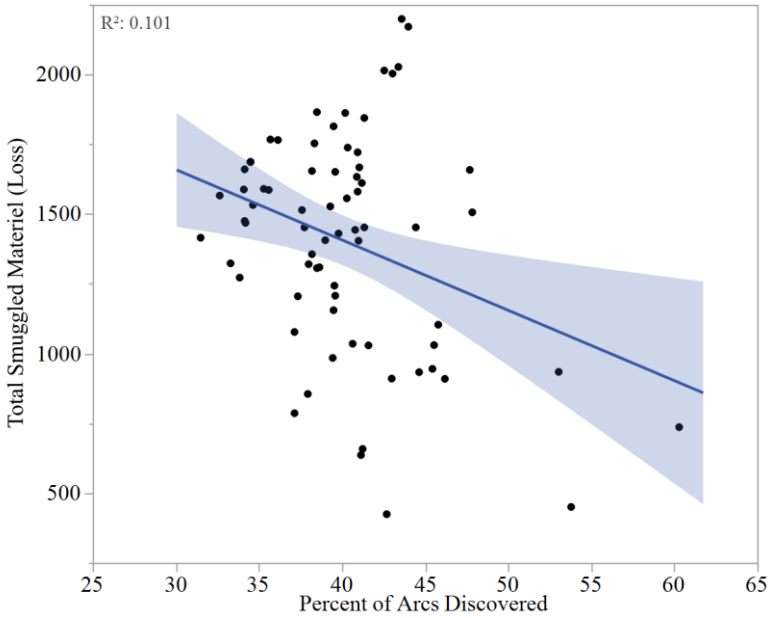

The relationship between the percentage of arcs discovered by the interdictor and the total smuggled materiel is very weak. The linear fit accounts for only 10 percent of the variance in smuggled materiel. The shaded region indicates the 95 percent confidence level.

(5) The Effect of Miscalculation

We examine the potential for the interdictor to miscalculate and the relationship of these miscalculations to performance. Owing to myriad indirect effects within this model, it is difficult to determine definitively if the interdictor has made a poor choice. To reduce confounded effects and test for miscalculation specifically, we add an experimental artifact to the network. As described in Section A.2, we design an arc, (*N13*, *N15*), into the Case 1 network that the interdictor must consider, but the smuggler can never use (Figure 48). The arc is an arc to nowhere. Even so, the interdictor is never made explicitly aware that smuggler transit is impossible on the arc. The interdictor must therefore consider the arc in his decision calculus and either address it as a threat or pay the arc little attention. If the interdictor expends resources against the arc, the expenditure would be a pure *miscalculation* of the smuggling threat. We find that the interdictor

miscalculates and places approximately 11 percent of the available sensors on arc (*N15, N13*) during 540 simulated games (Figure 49).

Figure 48. Arc to Nowhere.

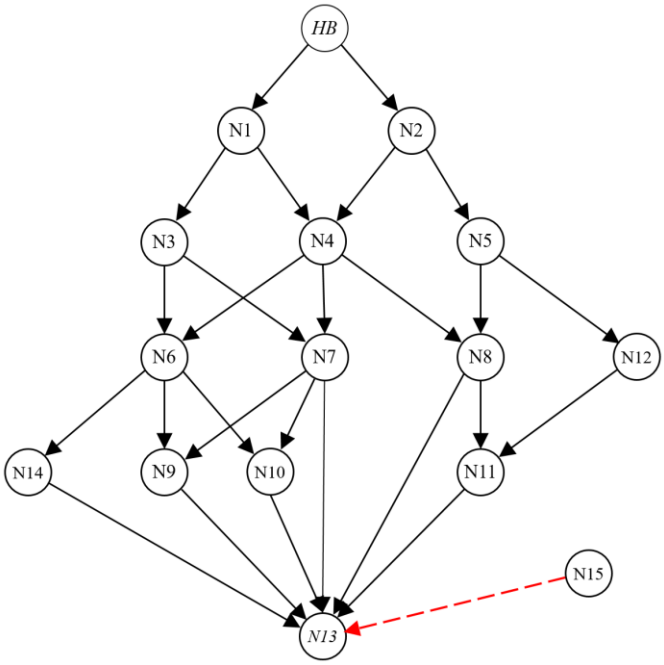

Arc (N15, N13) is part of no path available to the smuggler. Any placement of sensors on the arc by the interdictor would be a mistake, representing a *miscalculation*.

Figure 49. Percent of Sensor Budget Placed Inappropriately.

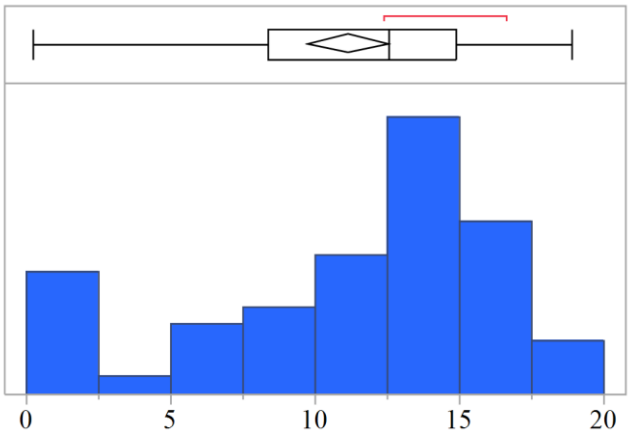

The distribution shows the percentages of the interdictor's available budget that he placed inappropriately on arc (*N15, N13*) during 540 simulated games. The mean percentage is 11% and the 95% Confidence Interval is [9.8%, 12.6%]. Note the non-symmetric shape and left-skewness.

Figure 50 shows the relationship between the total smuggled materiel and the percent of the interdictor's sensor budget placed inappropriately on arc (*N15*, *N13*). We partition the results by total sensor budget. At sensor budgets two and five, there appears to be a strong relationship between the amount of materiel smuggled and miscalculation. Both the total smuggled materiel and rate of miscalculation increase together. However, for sensor budgets three and four, there is a very weak relationship between the total materiel smuggled and percent of the interdictor's sensor budget placed inappropriately.

Figure 50. Total Smuggled Materiel versus Percent of Sensor Budget Misplaced by Sensor Budget.

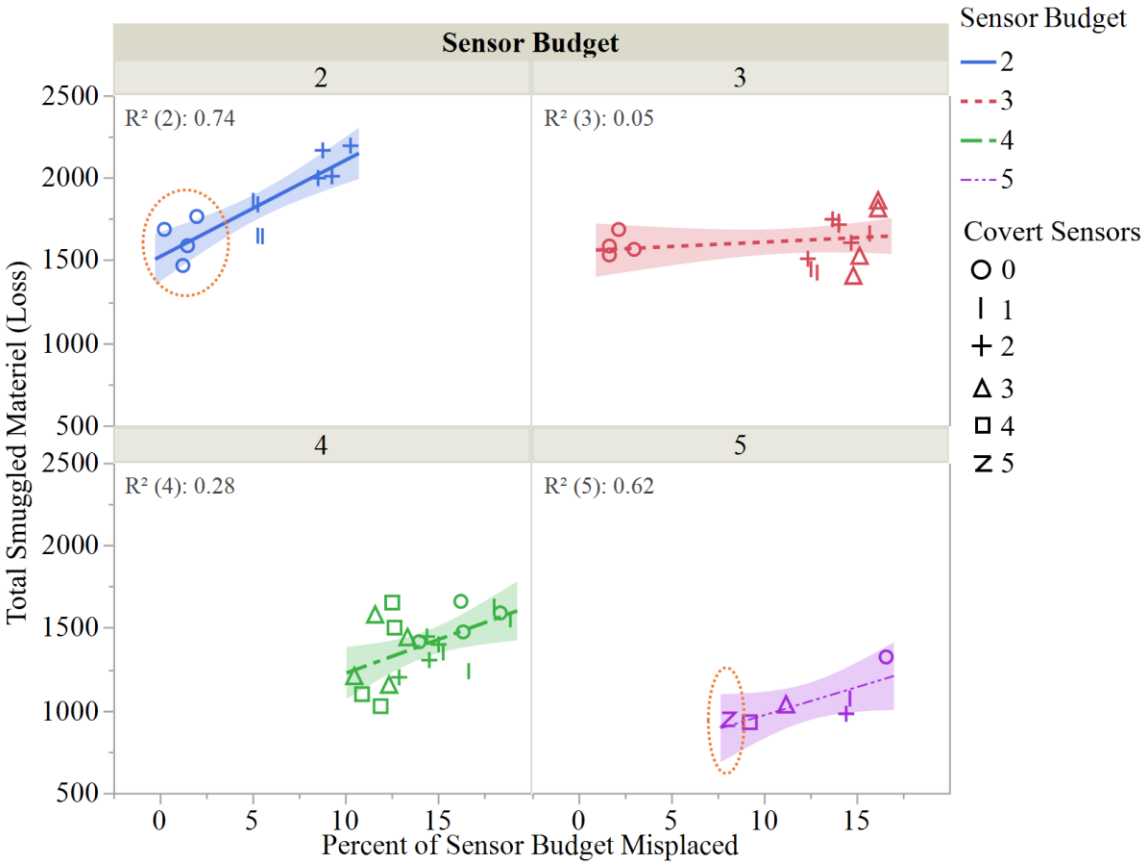

Upper Left and *Lower Right*: The relationship between the total smuggled materiel and rate of miscalculation is strongest for sensor budgets two and five, respectively. Lower rates of miscalculation correspond with lower total smuggled flow. However, the policy that results in the lowest rates changes from pure overt to pure covert (circled in orange). The shaded region indicates the 95 percent confidence interval for total smuggled materiel (loss).

We also find that as sensor budget increases, increasing the allocation to covert sensors decreases the miscalculation rate. With a budget of two sensors, a pure overt policy shows the lowest misclassification rates and total smuggler flow. However, as the sensor budget increases, policies with a higher proportion of covert sensors display superior performance in both misclassification rate and amount of smuggler network discovered. Thus, in spite of increased exploration, the miscalculation rate still falls significantly to very low rates. It is apparent that the interdictor develops a better "sense" of the network's topology with more covert sensors. We emphasize that the converse is not true; as shown above in Section (4), better knowledge of the topology does not always yield better performance in terms of minimizing the smuggled materiel.

Figure 51. Percent of Sensor Budget Misplaced versus Percent of Sensor Budget Allocated to Covert Sensors.

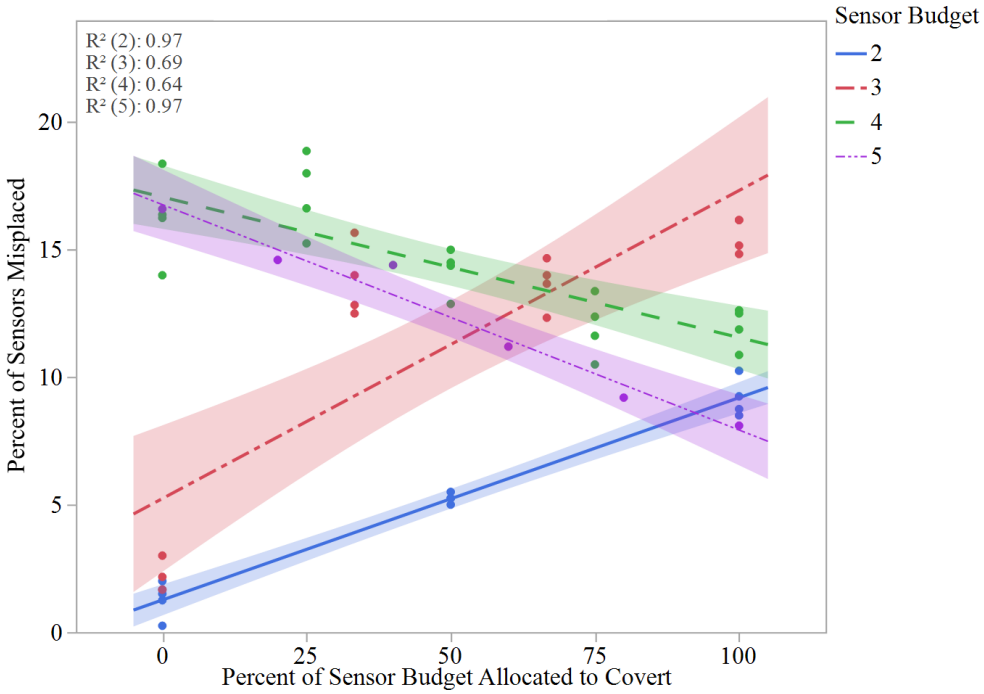

At sensor budgets two and three, policies that devote more budget to covert sensors miscalculate more often. With sensor budgets of four and five, the trend reverses. More covert sensors yields lower rates of miscalculation.

Figure 52. Percent of Arcs Discovered versus Percent of Sensor Budget Misplaced by Sensor Budget.

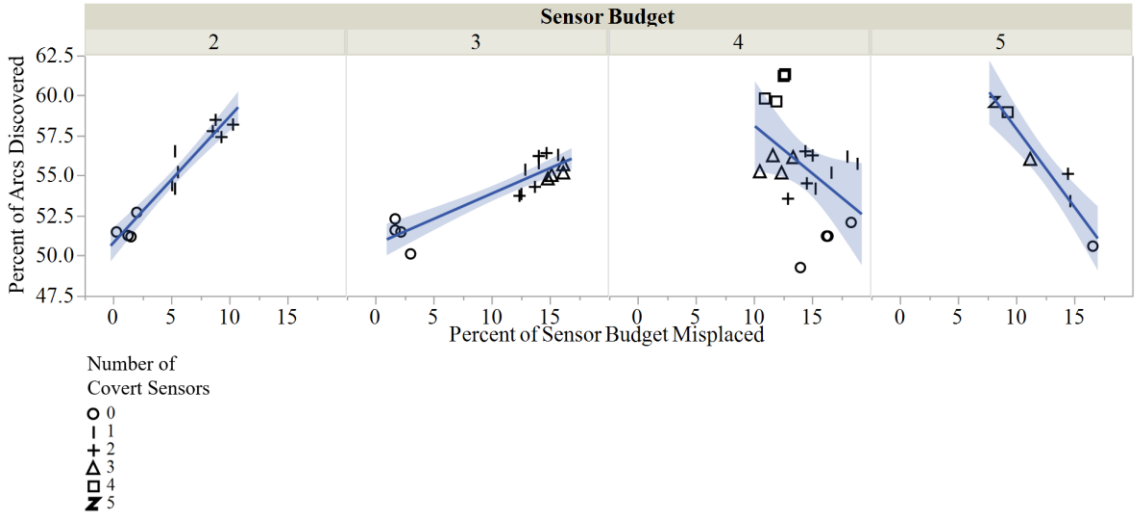

With increasing budget, increased allocation to covert sensors yields both more exploration and less miscalculation. *Left*: Budget with two sensors, hybrid and pure covert policies explore more but miscalculate more. *Middle Left*: Budget with three sensors, the trend line flattens, but hybrid and pure covert policies still miscalculate more. *Middle Right*: The trend reverses; pure covert and hybrid policies now explore more and miscalculate less. *Right*: The relationship is strong; covert-heavy policies explore significantly more and miscalculate significantly less.

3. Case 2

In Case 2, we consider a deeper and less wide smuggling network than that constructed in Case 1. There are 21 arcs connecting 14 nodes (Figure 53). The designed network of Case 2 provides 20 unique *s*-*t* paths from the node *HB* (*Hostile Base*) to node *N13*. Table 9 displays the initial network data. As in Case 1, the smuggler's estimates of arc capacities and the interdictor's estimates of arc costs and capacities all begin as 10, 10, and 1, respectively. We design the remaining network topology and data with the same motivation found in Case 1, to create a practical range of total route costs and reveal new routes over successive game rounds.

Figure 53. The Network for Case 2.

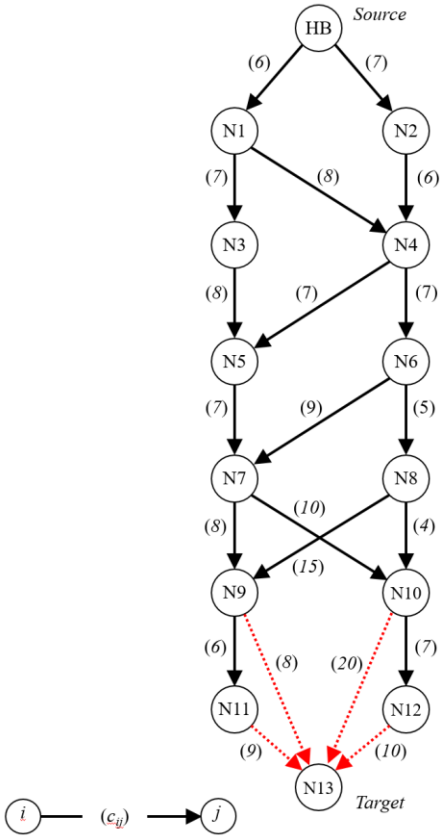

The network for Case 2 includes 14 nodes and 21 arcs. Arcs with cost above 10 are not visible. The interdictor could use four or more sensors to naively make an *s*-*t* cut across the reverse star of the target (dotted red).

In Case 2, we shorten the time horizon to 15 game rounds for all scenarios. Arcs with cost above 10 are not visible to the smuggler. The timer only introduces two arcs through the course of a game. We thus create a significantly different topology that is less fluid than that found in Case 1 to test further the extensibility of our observations.

Table 9. Case 2 Initial Interdictor and Smuggler Estimates.

Arc		Interdictor Estimates		Smuggler Estimates	
Tail	**Head**	**Cost**	**Capacity**	**Cost**	**Capacity**
N8	N9	10	1	15	10
N12	N13	10	1	10	10
N5	N7	10	1	7	10
N4	N6	10	1	7	10
N4	N5	10	1	7	10
N11	N13	10	1	9	10
HB	N2	10	1	7	10
N10	N13	10	1	20	10
N2	N4	10	1	6	10
HB	N1	10	1	6	10
N10	N12	10	1	7	10
N9	N11	10	1	6	10
N6	N7	10	1	9	10
N9	N13	10	1	8	10
N7	N10	10	1	10	10
N1	N3	10	1	7	10
N3	N5	10	1	8	10
N7	N9	10	1	8	10
N1	N4	10	1	8	10
N8	N10	10	1	4	10

In Case 2, we assume the smuggler's budget is 500 cost/round, just adequate to allow the smuggler to transport any materiel introduced at node *HB* in a single round through the network within the following two to three rounds. Some larger packets now require three rounds to traverse the network, owing to the greater network depth in Case 2. The master scheduling table introduces 186 total units of materiel at node *NB*. Under the same portioning scheme found in Case 1, the master scheduling table allots packets to each of the 15 rounds of play.

a. Design of Experiments

Similar to Case 1, we combine an Orthogonal Latin Hypercube (OLH), crossed design, and star points to consider 53 equally likely scenarios in Case 2. These scenarios also program two decision variables and five noise variables through a realistic range of values (Table 10). The 36 points of the augmented crossed design expand our analysis of scenarios with sensor budget between one and three. As in Case 1, in Case 2 we are

focused on situations where the interdictor cannot make a naïve *s-t* cut in the reverse star of the target (Figure 53). Four or more sensors would allow such a strategy. We use the experimental design to investigate the same three general policies found in Case 1: *pure overt*, *hybrid*, and *pure covert*. The design replicates each scenario at least 10 times to establish the minimum of 90 percent confidence of less than 10 percent error in the estimate of the expected value of total smuggled flow.

Table 10. Case 2 Range of Factors in Experimental Design.

Decision Variables	**Definition from Chapter 3**	**Range of values**
$covert.budget_t$	maximum number of covert sensors the interdictor can place in round *t* [cardinality/round]	[0, Sensor Budget]
$overt.budget_t$	maximum number of overt sensors the interdictor can place in round *t* [cardinality/round]	[0, Sensor Budget]
Noise Factors	**Definition from Chapter 3**	**Range of values**
$q_{(i,j),t}$	smuggler penalty for traversing an arc with sensor [cost/flow-round]	[5, 50]
$size_p$	size of packet *p* [flow]	[1, 2, 3, random integer [1,3]]
packet.signature	level of stealth for packets [non-negative integer]	[1, 2, 3]
attenuate	rate of cost attenuation $[\{attenuate \in \square \mid attenuate \in (0,1]\}]$	[0.88, 0.95]
γ	scaling parameter (0.0 – 1.0) that attenuates the interdictor's capacity estimate from the previous round, $\hat{u}_{(i,j),t-1}$	[0.6, 1.0]

b. Results

(1) Interdictor Policy Performance: Aggregate Metrics

Figure 54 shows the total smuggled flow versus the interdictor's sensor budget. As in Case 1, the plot for Case 2 clearly shows a face-valid trend of decreasing total flow with increasing sensor budget.

Figure 54. Total Smuggled Materiel versus Sensor Budget.

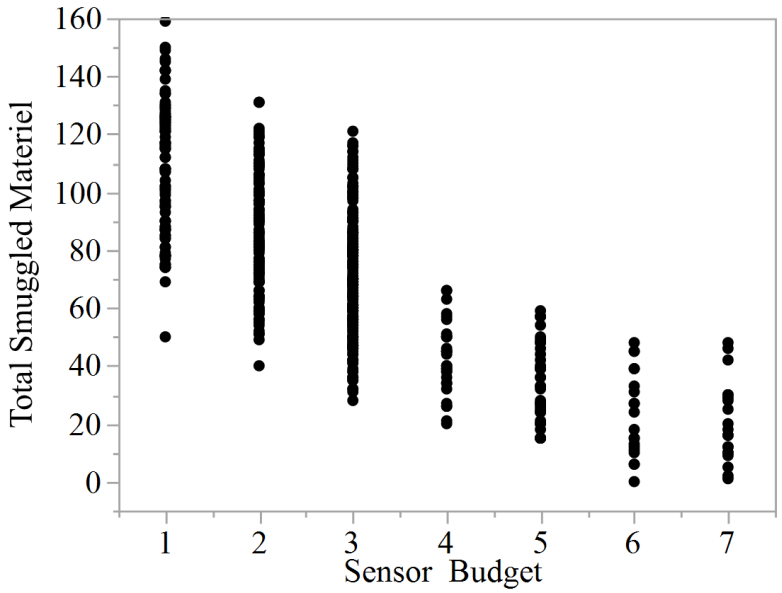

As in Case 1, the total smuggled flow per game in Case 2 decreases as the interdictor's sensor budget increases.

Figure 55 and Table 11 display the total degradation to smuggled flow and loss by policy. Figure 56 shows each individual policy's performance by sensor budget.

Table 11. Policy Performance: Percentage of Degradation of Smuggler Flow.

	Number of Covert Sensors			
Sensor Budget	0	1	2	3
1	***47%***	**36%**	NA	NA
2	***61%***	57%	**43%**	NA
3	68%	***70%***	64%	**49%**
Policy	***Pure Overt***	***Hybrid***	***Pure Covert***	

With a budget of one or two sensors, pure overt policies maximally degrade the smuggler flow. However, hybrid policies become more effective with a budget of three sensors, offering 70% degradation of the smuggler flow.

Figure 55. Policy Performance by Sensor Budget.

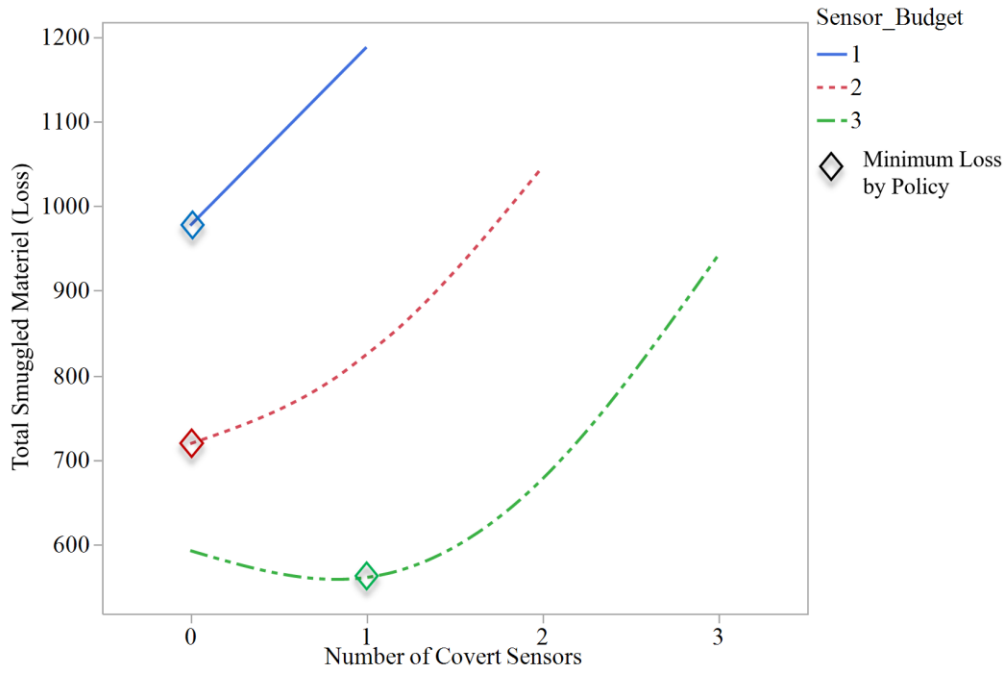

The projected loss contours show pure overt sensor policies as producing the minimum loss with a budget of two (blue line) and three sensors (red dashed). However, the hybrid policy is a superior allocation of three total sensors (green broken line). Model details are displayed in Appendix A.

Figure 56. Policy Performance by Sensor Budget.

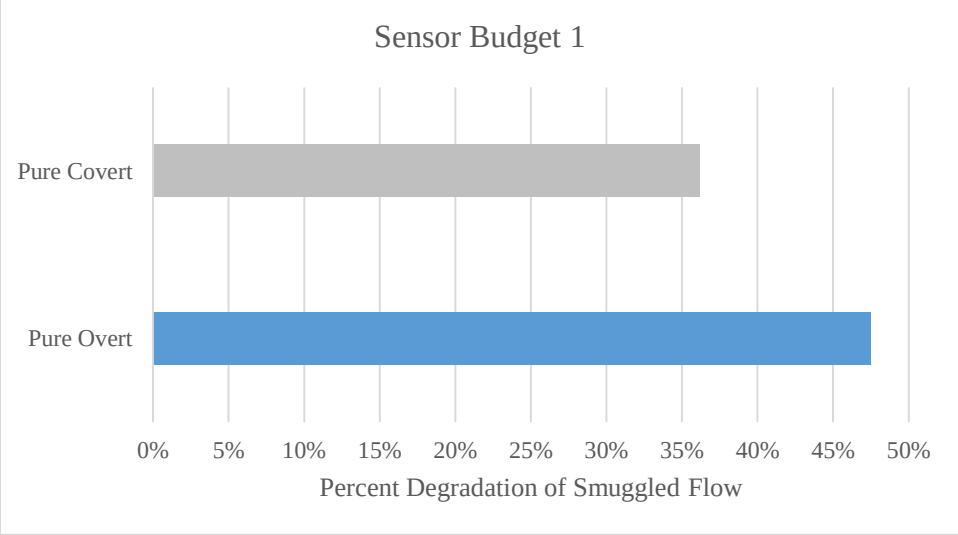

(a) With a sensor budget equal to one, the pure overt strategy maximally degrades the smuggled flow. It offers 9% more degradation than the pure covert policy.

Figure 56. Policy Performance by Sensor Budget (Cont.).

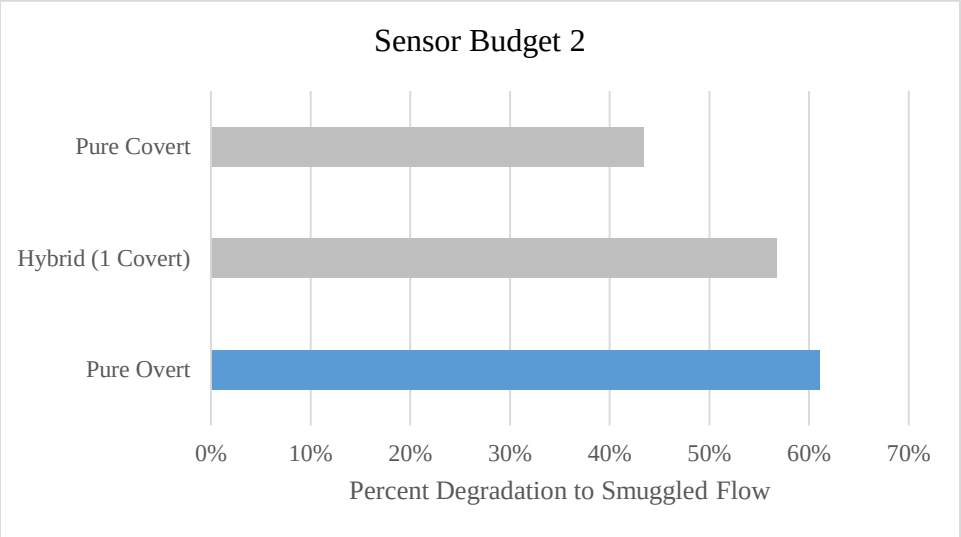

(b) The pure overt policy still offers the greatest degradation of smuggled flow with two sensors available. The degradation is 4% higher than the hybrid (1 covert) policy.

Figure 56. Policy Performance by Sensor Budget (Cont.).

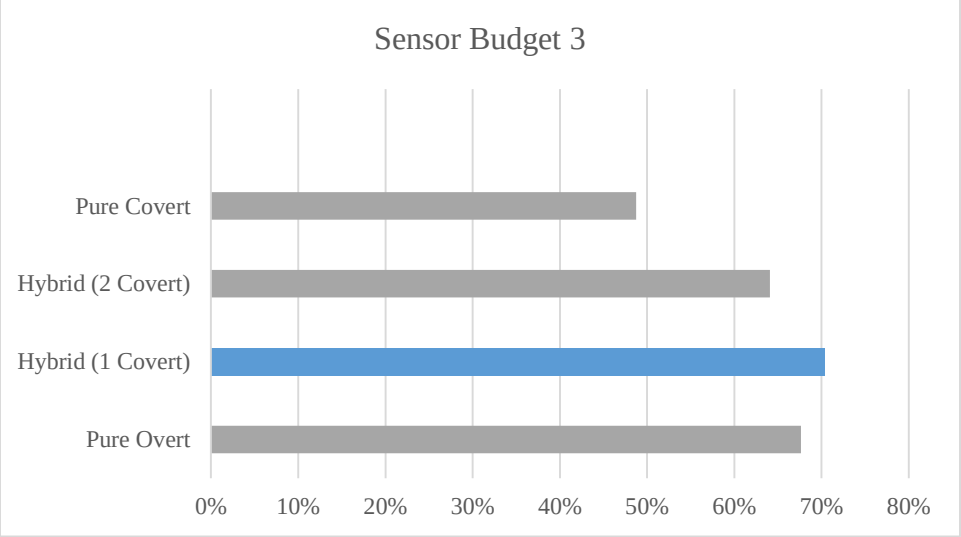

(c) With a budget of three sensors, we observe the hybrid (1 covert) policy cause maximum degradation of the smuggled flow.

As found in Case 1, increasing the probability of detection by an average of 8 percent yields increased degradation of the smuggled flow (Figure 57).

Figure 57. Effect of Sensor Sensitivity on the Percent Degradation of Smuggled Materiel.

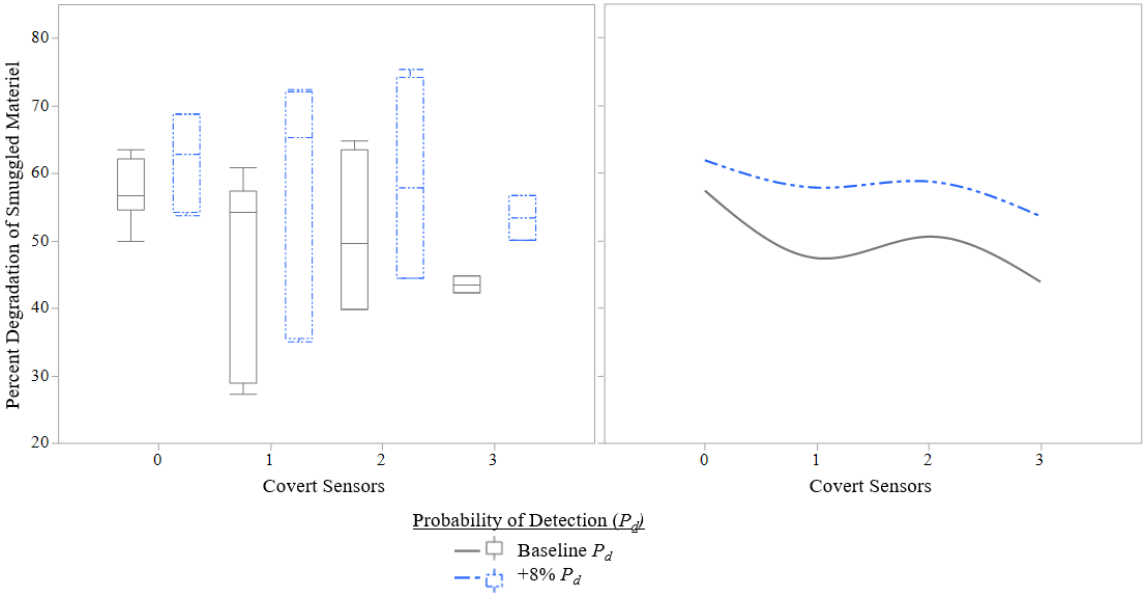

Similar to Case 1, improving each sensor's probability of detection by approximately 8% yields consistent improvement in all Case 2 scenarios. *Left*: Using a nonparametric permutation test, we can only distinguish between the two P_d scenarios with 80% confidence when the number of covert sensors is either one, two, or three. When there are three covert sensors, we can distinguish between the two P_d scenarios with over 99% confidence. *Right*: The average flow degradation increases from 5% to 11% as the number of covert sensors also increases. However, the largest increase in flow degradation, 15%, occurs at the *hybrid* (1 covert) policy. We produce the curves with piecewise splines, $\lambda = 0.05$.

The degradation generally increases as the interdictor enlists additional covert sensors, with a maximum 15 percent degradation of smuggled flow resultant from the *hybrid* (1 covert) policy.

(2) Interdictor Policy Performance: Time Dynamics

We continue our method of analysis established in Case 1 by proceeding to explore the time dynamics of each policy within Case 2. Using 530 individual time series, we construct the 0.9 quantile meta-game for each policy within each sensor budget (Figure 58).

In Case 2, we design a game that introduces only 2 arcs introduced by timer. The greater network stability manifests slightly more stable game play. While still oscillating, the smuggled flow does not change round-to-round by almost an order of magnitude as observed in Round 8 of Case 1. Even so, at sensor budget 3, the hybrid policy (2 covert) offers the same buffering of peak round-to-round flow we saw in Case 1.

Figure 58. *Meta-Games*, the 0.9 Quantile of Smuggled Flow by Game Round.

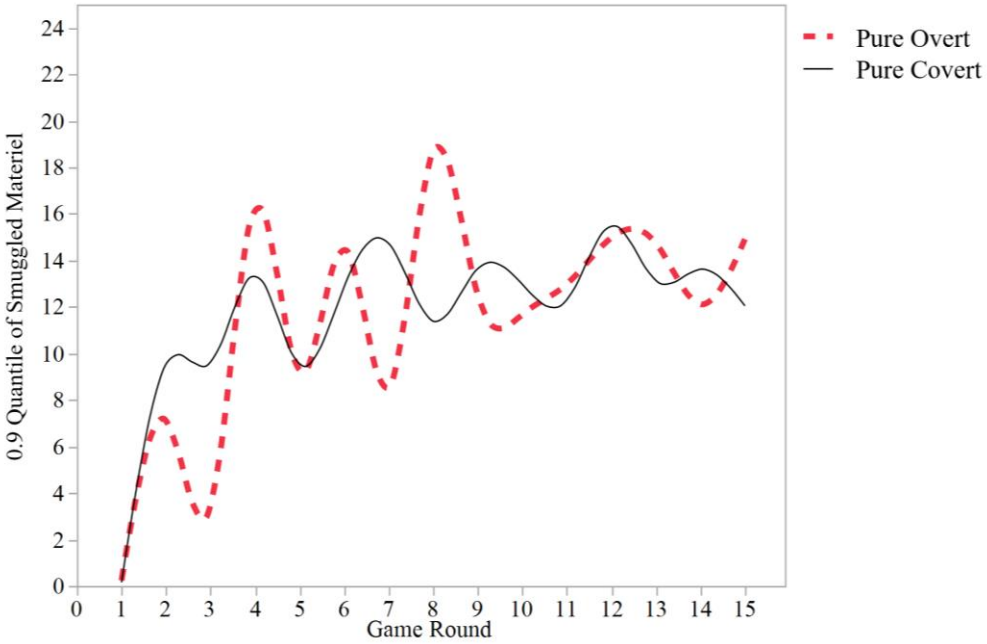

(a) **Sensor Budget 1**. Both the peaks and troughs of round-to-round flow are more extreme under the pure overt policy.

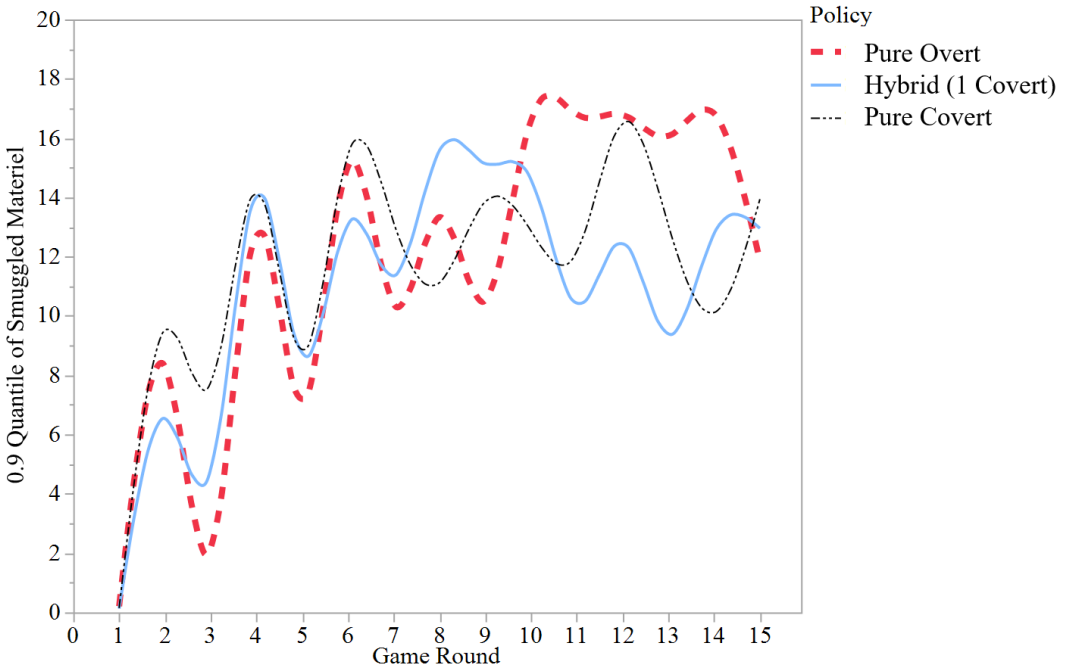

(b) **Sensor Budget 2**. There is no significant difference between each policies' round-to-round performance with only two sensors available. Policies performing well in some rounds perform poorly elsewhere.

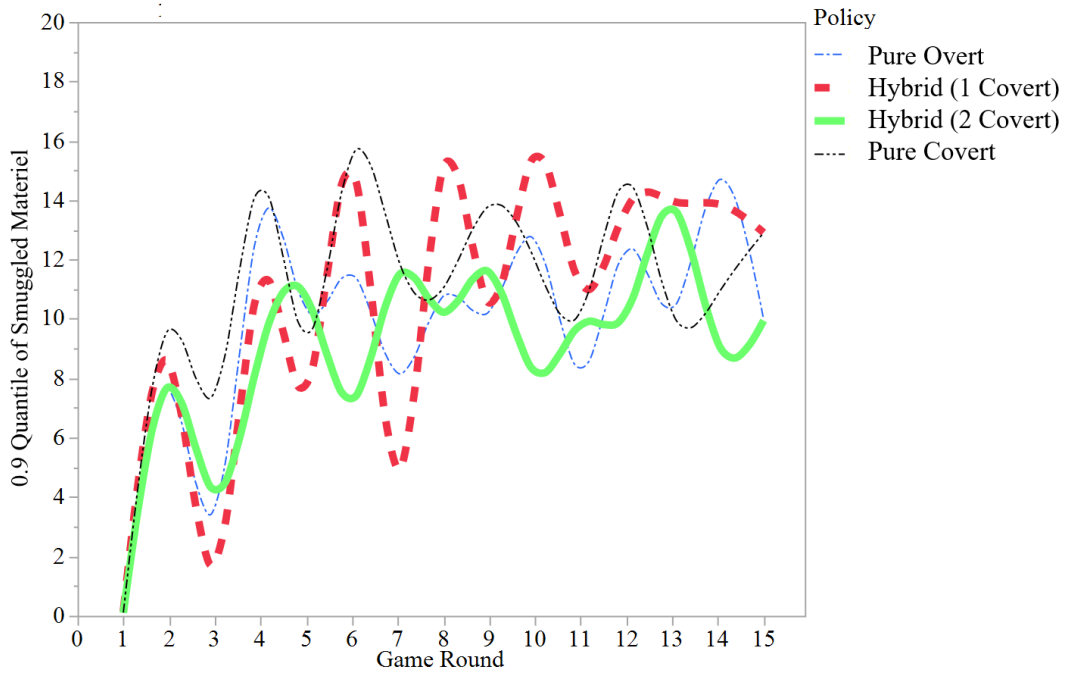

(c) **Sensor Budget 3**. *The policies now distinguish themselves*. We note the hybrid (2 covert) policy permits approximately 30–40 percent less round-to-round peak smuggled flow than the hybrid (1 covert) policy.

(3) Evaluation of Seizures as Proxy for Smuggled Flow

In Case 2, the value of inference achieved when comparing the total materiel detected to the total smuggled is highly inconsistent. Even with a variety of transformations of the independent variable, the total materiel seized by the interdictor, linear regression models continue to indicate both heteroscedasticity and weak explanatory power (Figure 59). Consistent inference is only available once we partition the results by individual policy (Figure 60). At sensor budget 1 and 2, the variation in the total materiel detected explains between 72 and 91 percent of the variation in total smuggled flow. As in Case 1, the strongest relationships occur when the interdictor assigns a large proportion of his sensor budget to covert sensors. The amount of detected materiel is very weakly related to the total smuggled flow for pure overt and hybrid (2 covert) (Figure 61). The R^2 for the best corresponding linear models is 0.07 and 0.50, respectively.

Figure 59. Linear Regression of Total Smuggled Materiel versus Total Materiel Detected.

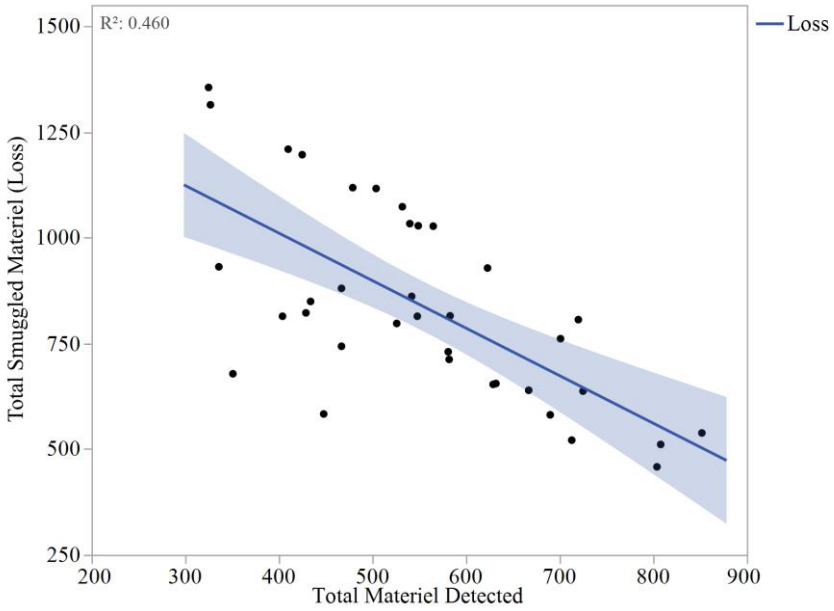

A linear model between total smuggled materiel and total materiel destroyed suffers from both heteroscedasticity and moderate explanatory power.

Figure 60. Total Smuggled Materiel versus Total Materiel Detected.

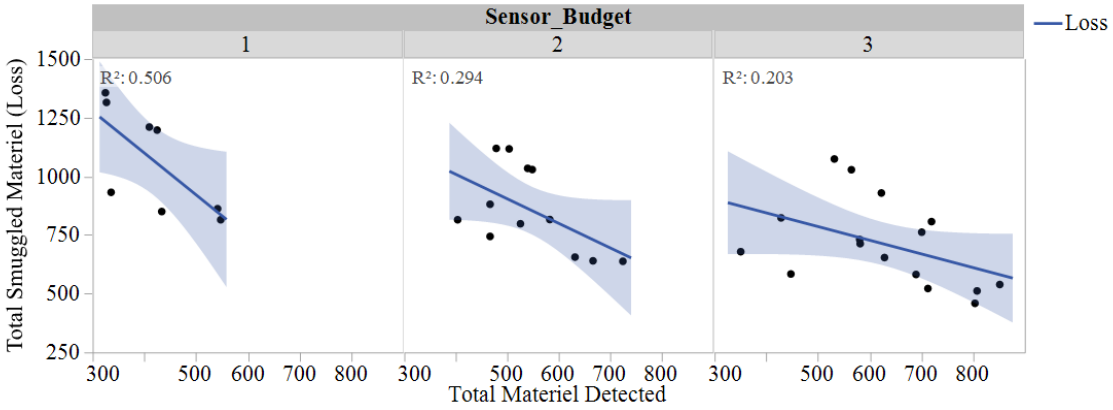

Naïve linear models show extremely weak relationships between the total smuggled materiel and total materiel detected by the interdictor.

Figure 61. Model of Total Smuggled Materiel versus Total Materiel Detected by Policy at Sensor Budget 2.

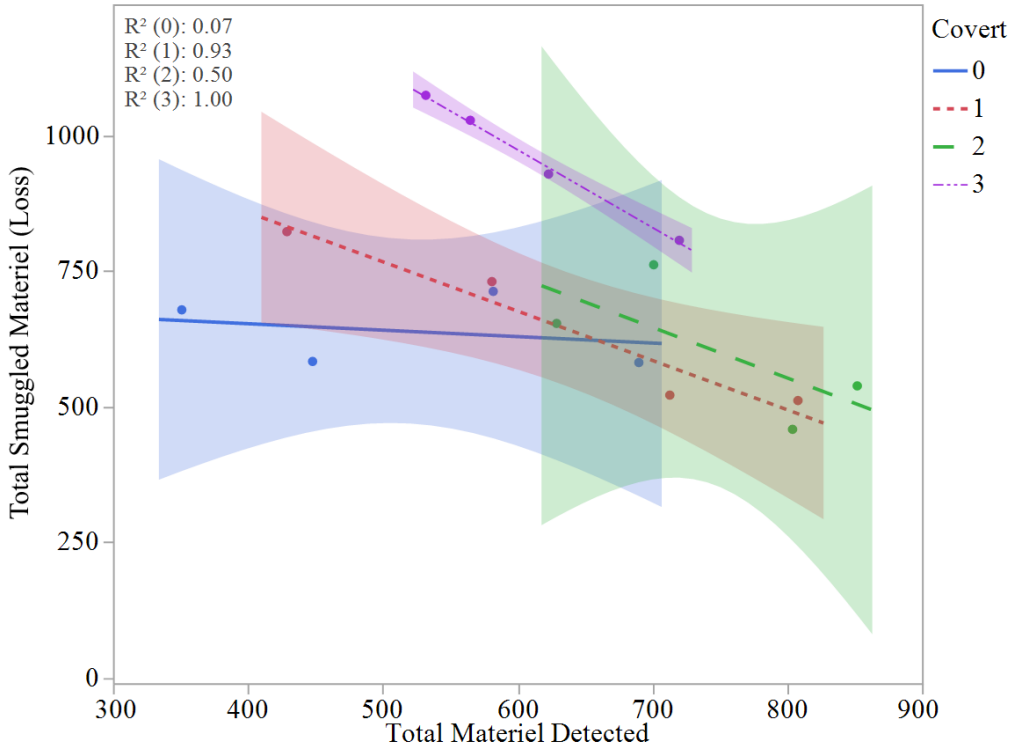

As with Case 1, with a budget of four sensors, the relationship between the total materiel detected and the total smuggled materiel is highly inconsistent. It varies from high explanatory power for pure covert policy to very low explanatory power for pure overt and hybrid (2 covert). The shaded regions represent a 95 percent confidence interval.

(4) The Relationship between Discovered Network and Flow

We again explore the relationship between the interdictor's success at degrading the smuggled flow and the scope of the smuggling network the interdictor discovers. Our Case 2 results are consistent with those from Case 1; there is a very weak relationship between the total smuggled flow and the percentage of the network the interdictor discovers. The relationship between the proportion of covert sensors within the interdictor's policy and the percent of the network he discovers is also non-monotonic and does not fully capture the variance in the smuggled flow (Figure 62). Using a non-parametric scaled Kruskal-Wallis H test, there is strong evidence that as the allocation of covert sensors increases, so does the percent of arcs discovered change (p-value ~ 10^{-6}).

We continue assessing the value of information by now quantifying it with the total amount of materiel smuggled. Figure 63 shows the total smuggled materiel versus the percentage of all arcs discovered by the interdictor within each of the 47 Case 2 scenarios.

Figure 62. Percent Arcs Discovered versus Percent of Sensor Budget dedicated to Covert Sensors.

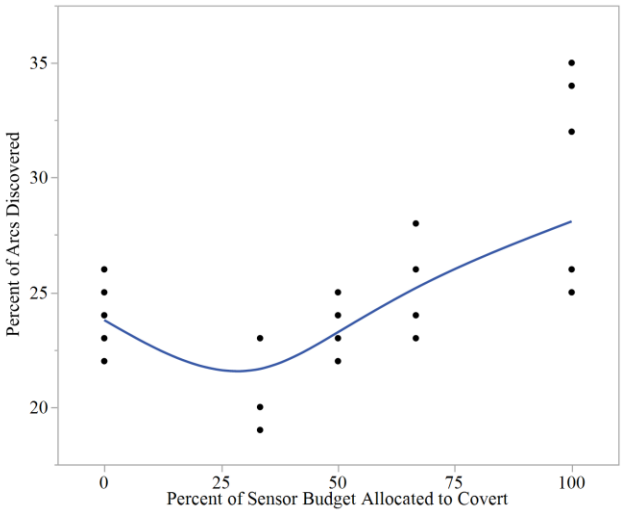

After one covert sensor is included, policies with more covert sensors offer increasing percentages of discovered network. The results of the policies are statistically distinguishable (p-value: 10^{-6}). The curve represents a nonparametric kernel density smoother model using local weights (blue line).

As in Case 1, we fit a number of statistical models in an attempt to explain the amount of smuggled flow by the percent of arcs discovered. However, the best fit shows an extremely weak relationship, accounting for only 16 percent of the variation in total smuggled flow ($R^2 = 0.161$).

Figure 63. Percentage of Arcs Discovered versus Total Loss.

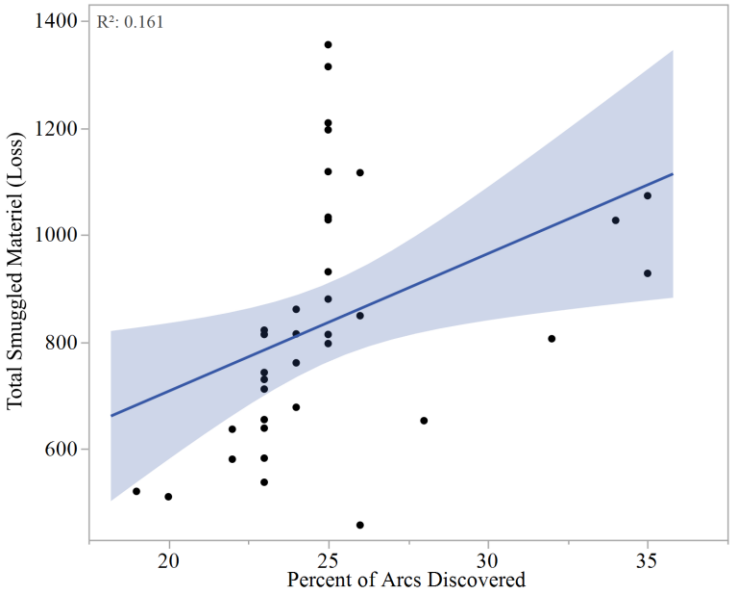

The relationship between the percentage of arcs discovered by the interdictor and the total smuggled materiel is very weak. The best of a variety of model fits only explains 16 percent of the variance in smuggled materiel. The shaded region indicates the 95 percent confidence level.

Using a linear model, the percentage of arcs discovered only explains 10 percent of the variance in the total amount of materiel smuggled. Further investigation by specific interdictor policy also revealed weak relationships between these two factors.

B. OBSERVATIONS AND DISCUSSION

Our two computational cases provide consistent results and expose several important unanticipated insights. Based on the data provided by these two cases, we now draw important inference and expose several relationships.

When the budget for sensors is small, relative to the number of smuggling routes, the interdictor should allocate his entire budget to overt sensors. Overt policies tend to capture fewer smuggled packets, but raise the smuggler's perceived risk significantly. Observing the placement of multiple overt obstacles, the smuggler shows risk aversion and smuggled flow rapidly becomes frustrated within the route structure. The emergent risk-averse behavior is analogous to deterrence. Using our model, we are able to assess the value of the deterrence, given other policy options.

As the budget for sensors increases, it becomes more effective for the interdictor to employ a hybrid policy involving a mix of overt and covert sensors. In effect, the interdictor uses the smuggler's perception of overt sensors to increase the effectiveness of covert sensors. The interdictor can employ hybrid policies in an ambush pattern, herding smuggled flow with overt sensors into covert sensors that lay in wait. That is, addressing a large front, the risk of the overt sensors substantially increases the attractiveness of the route on which the interdictor placed a covert sensor, often making the route the smuggler's new primary choice. Generally, the interdictor has to be less accurate in his estimations to employ hybrid policies effectively.

In contrast to overt policies, pure covert policies primarily achieve their effect by actually seizing the smuggled flow. Deterrence is far less important under pure covert schemes. For these policies, information is key. Pure covert policies encourage the interdictor to target arcs deeper within the smuggling network more precisely. The degree of exploration is highest under these policies. Unfortunately, the increased degree of exploration does not also increase the interdictor's effectiveness. With a dynamic topology and extremely agile smuggler, the interdictor is often precisely wrong. Orienting on a specific set of paths, the interdictor is drawn further into the smuggling network, only to be quickly out-flanked by the smuggler's reaction. The interdictor thus fairs poorly in instances of low sensor budgets and pure covert policies, unable to form consistently accurate estimates of the smuggler's intentions.

Counter to our intuition, our examination of each policy's time dynamics reveals that policies that most-reduce the total smuggler flow are not always the same policies that perform best round-to-round. We expect policies generating the highest degradation

to total smuggled flow to have also the smallest maximum value of flow in any single game round. Our expectation is confirmed in instances of low interdictor sensor budgets. However, in both Case 1 and Case 2, instances of higher sensor budgets reveal that hybrid policies with more covert sensors reduce the peak round-to-round flow considerably.

Hybrid policies appear to prevent surprise to a substantially greater degree. Case 1, round 8 provides the tactical analogy of smuggler deception and well illustrates the improved round-to-round performance of hybrid policies described above. As a new route is withheld until round 8, the interdictor is seduced to varying degree by the smuggler's—albeit unconscious—deceptive actions in rounds 1 to 7. The interdictor policies that are least deceived in rounds 1 to 7 are not always those that most reduce the total smuggled flow. However, in almost all cases we explored they are a close second.

It appears that while adaptability is important in the interdictor-smuggler problem, under situations that consider two-sided uncertainty, being too agile and too aggressive in exploration can significantly reduce effectiveness. Under several pure covert policies, the interdictor's feedback loops appeared overly sensitive, causing frequent and inappropriate changes to sensor locations round-to-round. Smuggler flow was distinctly higher when this happened. It appears that the adventurous, under-informed, and hyper-agile interdictor fluctuates wildly. Hybrid policies appear to calm the fluctuation, and appear to leverage the smuggler's own risk calculus against him.

Presented with the opportunity to increase or develop his resources, the interdictor must prioritize his goals. With minimal budgets, the interdictor should focus on maximizing deterrence. Gaining an additional sensor is more helpful in these cases than increasing sensor sensitivity. With greater budgets, establishing a hybrid policy should be the interdictor's priority. Once established, increasing covert sensor sensitivity provides significant rewards in reduced total smuggled flow, higher seizure rates, and greater resilience to surprise. We have not modelled the effect of intelligence gained by exploiting seized materiel. The knowledge from these exploits can prove tactically decisive. More heavily covert hybrid policies offer substantial increases in materiel capture rates and should be selected if materiel exploitation is a priority.

We expect that as the interdictor discovers more of the smuggler's network, the interdictor would be able to better influence the smuggler flow, but this is not the case. The observation reinforces our earlier warning on the dangers of interdictor adventurism under limited information. Even so, if discovery of the physical structure of the network is important, policies with heavy covert sensor allocations are best suited to the task.

Our limited excursion to study the effect of miscalculation demonstrates that the policies that best reduce total smuggled flow and are more resilient to surprise also result in the least miscalculation by the interdictor. This finding corroborates the policy performance assessments thus far. As before, the reduced miscalculation is generally not related only to the type of policy, but instead related to the type of policy in light of the sensor budget.

Lastly, we show that the amount of seized materiel is a poor proxy for the total amount of smuggled flow. The relationship between these factors is very inconsistent. In policies where deterrence is high, the amount of seized materiel provides almost no information on the actual amount of unseen materiel successfully smuggled. Now presented with the results, the finding agrees with our intuition. A policy that relies on deterrence, an intangible emergent effect, should not use seizures as a performance measure. In contrast, the amount of materiel seized can be used to very accurately estimate the total smuggled flow for heavily covert policies. We believe it is impractical to distinguish these situations in real world interdictor-smuggler contests. Estimates that naïvely measure overall interdiction performance by seizure rate should be carefully examined.

V. CONCLUSION

The main points of this thesis are twofold:

(1) A range of realistic, complex behaviors can emerge from the interaction of two hostile, intelligent agents acting under simple rules within a dynamic environment of uncertainty and danger.

(2) We can gain insight into these complex situations through heuristics that combine complimentary optimization, stochastic, and game-theoretic models under an umbrella of simulation.

Previous research into martial contests, such as this interdictor-smuggler context, has made various strong assumptions in the name of tractability. These persistent assumptions include perfect information, an unchanging environment, and non-adaptability. Our constructive cases from Chapter 3 and full model results in Chapter 4, demonstrate massive perturbations in both the conduct of play and game results as we deliberately and cumulatively relax these assumptions. Each of these relaxations admits an essential feature of the problem—features unambiguously identified by the breadth of military doctrine as the essence of conflict. The significant effect of these essential features, demonstrated in this thesis, should draw sharp focus on the results of many models that attempt to study combat otherwise. Specifically, is it sensible to study situations in a martial context and ignore uncertainty or adaptation? Assigning omnipotence to one or both combatants in this type of context for tractability is deeply dissatisfying. It conflicts at a philosophical level with both historical and contemporary military thought. Solving problems under a philosophically different paradigm than that in martial practice endangers the real utility and accuracy of any insights gained. Solving problems by artificially striping away essential complexity so that it fits more cleanly into only optimization, game theory, or stochastics is equally limiting.

We have shown that constructing an interdisciplinary model can unite the strengths of these fields in a complimentary fashion. It can generate feasible face-valid solutions to our problem at very limited computational expense. By also considering time

dynamics, many new insights present themselves, such as the emergence of deterrence, giving light to new types of research questions.

Our study is meant to be a prototype, demonstrating the power of a hybrid model. It is not without its limitations. We have made assumptions on the method and speed by which hostile agents might evolve in a martial context. While these provide a direct analogue to contemporary military methods, there are other ways. We did not include the advantage obtained by interrogation and exploitation after the capture of enemy forces and materiel. In many military campaigns, intelligence gained through such exploitation has been decisive. The duration of our cases considered was necessarily finite. Even so, the basis of our time horizon is subject to debate; so is the design and programmed evolution of our network topologies.

The degree of validation for this type of research may be significantly bounded by the limits of *knowability* inherent to these problems. These limits have to do with the observability and measurability of important performance data within the high uncertainty and mortal danger of real-world combat. In many ways, refinements of computational simulation using a blend of optimization, stochastics, and game theory may be the best way to continue investigation. Even so, future research could admit real world data to craft a wider array of smuggling networks or specific instances of interest. Under a broader set of configurations, models similar to ours could prove a great aid to training both inter-agency decision makers and their staffs. That training could encourage unique perspectives and seed important questions that might expose highly non-intuitive and indirect ways of influencing the outcome and assessing performance during real interdiction or counter-trafficking missions.

Further investigation should also more deeply consider the coupling of an interactive environment and intelligent agents. By allowing this interaction, we admit a difficult yet realistic miasma of uncertainty, risk, estimation, prediction, miscalculation, success, failure, deception, and deterrence. Within such a challenging scene, unreconciled differences between orientation and reality can bring catastrophe, despite adequate resources. Policies that encourage deliberate and beneficial adjustments of orientation while allowing timely action in spite of uncertainty beg greater study. The value of

relative information superiority is likely involved. Indeed, on a deeper level, there may be far stronger links between perception, performance, adaptation, innovation, and the rules by which we both consciously and unconsciously choose to operate.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX. STATISTICAL MODELS

Case 1: General Additive Regression Model

	***Dependent Variable**: Total Smuggled Materiel (Loss)*		
Term	**Coefficient**	**Std Error**	**Prob>|t|**
Intercept	2229.8197	70.56683	<0.0001
Covert	-152.1365	13.06025	<0.0001
Overt	-206.9444	13.52863	<0.0001
Q	-0.654293	1.119532	0.5614
Size{2&3&1-4}	-113.0372	23.80764	<0.0001
Size{2-3&1}	18.667449	29.5781	0.5306
Size{3-1}	-0.589296	37.26125	0.9874
Signature	-125.62	12.49998	<0.0001
(Covert-1.98611)*(Overt-1.98611)	-22.24899	8.598785	0.0124
(Covert-1.98611)*Size{2&3&1-4}	105.55112	17.07788	<0.0001
(Covert-1.98611)*(Signature-2.09722)	-20.83159	8.214892	0.0141
(Overt-1.98611)*Size{2&3&1-4}	30.366644	14.54132	0.0415
(Overt-1.98611)*Size{3-1}	-123.7667	27.57538	<0.0001
(Q-30.0556)*Size{2-3&1}	-21.87887	3.4043	<0.0001
(Q-30.0556)*Size{3-1}	-7.854226	3.561776	0.0317
Size{2-3&1}*(Signature-2.09722)	419.73442	71.7034	<0.0001
(Covert-1.98611)*(Covert-1.98611)	-36.33692	8.211495	<0.0001
(Overt-1.98611)*(Overt-1.98611)	36.884488	5.599369	<0.0001
R^2	0.96		
Adjusted R^2	0.95		

Case 2: General Additive Regression Model

	***Dependent Variable**: Total Smuggled Materiel (Loss)*		
Term	**Estimate**	**Std Error**	**Prob>|t|**
Intercept	1315.2894	39.21573	<0.0001
Covert	-125.6307	11.74715	<0.0001
Overt	-252.6989	12.73089	<0.0001
Signature	-62.22282	14.45549	<0.0001
(Overt-1.43396)*(Overt-1.43396)	60.392199	9.516845	<0.0001
R2	0.90		
Adjusted *R2*	0.90		

LIST OF REFERENCES

Alderson D, Brown G, Carlyle M, Cox L (2013) Sometimes there is no ‘‘most-Vital’’ Arc: Assessing and improving the operational resilience of systems. *Military Operations Research.* 18(1): 21–37.

Almasy S, Meilhan P, Bittermann J (2015) Paris massacre: At least 128 killed in gunfire and blasts, French officials say. *CNN Online* (November 14), http://www.cnn.com/2015/11/13/world/paris-shooting/.

Auer P, Cesa-Bianchi N, Freund Y, Schapire R (1995) Gambling in a rigged casino: the adversarial multi-armed bandit problem. *Proceedings, 36th Annual Symposium on Foundations of Computer Science* (Institute of Electrical and Electronics Engineers, Washington, DC), 322–331.

Awerbuch B, Kleinberg R (2004) Adaptive routing with end-to-end feedback: distributed learning and geometric approaches. *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing* (ACM, Chicago, IL), 45–53.

Banco E (2015) Supply and demand: European smugglers are the winners as new routes open up for refugees fleeing conflict in the Middle East. *International Business Times Online* (October 6), http://www.ibtimes.com/supply-demand-european-smugglers-are-winners-new-routes-open-refugees-fleeing-2127968.

Borrero J, Prokopyev O, Sauré D (2015) Sequential shortest path interdiction with incomplete information. *Decision Analysis Online* (November 19), http://dx.doi.org/10.1287/deca.2015.0325

Boyd, John R (1996) The Essence of winning and losing. Unpublished Lecture Notes from 1976 briefing, https://fasttransients.files.wordpress.com/2010/03/essence_of_winning_losing.pdf

Bubeck S (2011) Introduction to online optimization. Unpublished lecture notes, Princeton University, Princeton.

Cesa-Bianchi N, Lugosi G (2006) *Prediction, Learning, and Games* (Cambridge University Press, New York).

Cesa-Bianchi N, Lugosi G (2012) Combinatorial bandits. *Journal of Computer and System Sciences.* 78(5): 1404–1422.

Corley H, Sha D (1982) Most vital links and nodes in weighted networks. *Operations Research Letters.* 1(4): 157–160.

Cormican K, Morton D, Wood R (1998) Stochastic network interdiction. *Operations Research.* 46(2): 184–197.

Command, U.S. Naval Doctrine (1995) *Naval Doctrine Publication 6: Naval Command and Control* (Norfolk, VA).

Danskin J (1966) The theory of max-min, with applications. *SIAM Journal on Applied Mathematics.* 14(4): 641–664.

Dilanian K (2016) U.S. says it's slowing flow, but foreign fighters still flock to ISIS. *NBC News Online* (January 16), http://www.nbcnews.com/news/world/u-s-says-it-s-slowing-flow-foreign-fighters-still-n494281.

Dudewicz E, Dalal S (1975) Allocation of observations in ranking and selection with unequal variances. *Sankhya*. B37: 28–78.

Elder G (2006*)* Intelligence in war: it can be decisive. *Studies in Intelligence*. 50(2): 13–25.

Fulkerson D, Harding G (1977) Maximizing the minimum source-sink path subject to a budget constraint. *Mathematical Programming.* 13(1): 116–118.

GAMS Development Corporation (2013) General Algebraic Modeling System (GAMS) Release 24.2.1 (Washington, DC).

Ghare P, Montgomery D, Turner W (1971) Optimal interdiction policy for a flow network. *Naval Research Logistics Quarterly*. 18(1): 37–45.

Hagberg A, Schult D, Swart P (2008) Exploring network structure, dynamics, and function using NetworkX. Varoquaux G, Vaught T, Millman J, eds. *Proceedings of the 7th Python in Science Conference (SciPy2008)*. (Pasadena, CA USA), 11–15.

Harris T, Ross F (1955) Fundamentals of a method for evaluating rail net capacities. *Research Memorandum RM-1573* (The RAND Corporation, Santa Monica, California).

Headquarters, Department of the Army (2012) *ADP 6–0: Mission Command* (Headquarters, Department of the Army, Washington, DC).

Headquarters, Marine Corps Combat Development Command (1997) *Marine Corps Doctrinal Publication 1: Warfighting* (United States Marine Corps, Washington, DC).

Headquarters, Marine Corps Combat Development Command (2001) *Marine Corps Doctrinal Publication 1–0: Marine Corps Operations* (Department of the Navy, Washington, DC).

Israeli E, Wood R (2002) Shortest-path network interdiction. *Networks*. 40: 97–111.

JMP Pro (2015) Version 12.0.1 (SAS Institute Inc., Cary, NC).

Kleijnen, J, Sanchez S, Lucas T, Cioppa T (2005) A user's guide to the brave new world of designing simulation experiments, *INFORMS Journal on Computing*, 17(3): 263–289.

Kirby M (2003) *Operational Research in War and Peace: The British Experience from the 1930s to 1970* (Imperial College Press, London, UK).

Lehman J, Phelps S (2004) *West's Encyclopedia of American Law-Po-San* (Thomson & Gale, Detroit).

Malik K, Mittal A, Gupta S (1989) The k most vital arcs in the shortest path problem. *Operations Research Letters*. 8(4): 223–227.

Matsumoto M, Nishimura T (1998) Mersenne Twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 8(1): 3–30.

Morton D, Pan F, Saeger K (2007) Models for nuclear smuggling interdiction. *IIE Transactions.* 39(1): 3–14.

Nehme M (2009) Two-person games for stochastic network interdiction: models, methods, and complexities. Unpublished doctoral dissertation, The University of Texas, Austin.

Omrani B (2009) The Durand Line: history and problems of the Afghan-Pakistan border. *Asian Affairs* 40(2): 177–195.

Prados J (1999) *The Blood Road: The Ho Chi Minh Trail and the Vietnam War* (John Wiley & Sons, New York).

R Core Team (2016) *R: A language and environment for statistical computing* (R Foundation for Statistical Computing, Vienna, Austria).

RAND Corporation. History and mission. Accessed February 15, 2016, http://www.rand.org/about/history.html.

Ratliff D, Sicilia G, Lubore S (1975) Finding the n most vital links in flow networks. *Management Science.* 21(5): 531–539.

Robbins H (1952) Some aspects of the sequential design of experiments. *Bulletin American Mathematical Society.* 55: 527–535.

Rubin, Barnett R (2000) The political economy of war and peace in Afghanistan. *World Development* 28(10): 1789–1803.

Salmeron J (2012) Deception tactics for network interdiction: a multiobjective approach. *Networks.* 60(1): 45–58.

Sanchez S (2000) Robust Design: Seeking the best of all possible worlds. *Proceedings of the 2000 Winter Simulation Conference* (Institute of Electrical and Electronic Engineers, Piscataway, NJ), 69–76.

Sanchez, S. M. 2011. *NOLHdesigns spreadsheet*. Accessed November 11, 2015, http://harvest.nps.edu/.

Sharma, S (2015) Map: how the flow of foreign fighters to Iraq and Syria has surged since October. *Washington Post* (January 27), https://www.washingtonpost.com/news/worldviews/wp/2015/01/27/map-how-the-flow-of-foreign-fighters-to-iraq-and-syria-has-surged-since-october/

Shoichet, C (2016) Brussels attacks: charges filed, a man freed and suspects on the run. *CNN Online* (March 28), http://www.cnn.com/2016/03/28/europe/brussels-investigation/index.html.

Stewart, Van den Honert, eds. (2013) *Trends in Multicriteria Decision Making: Proceedings of the 13th International Conference on Multiple Criteria Decision Making*. (Springer Science & Business Media, Berlin).

Tausworthe R (1965) Random numbers generated by linear recurrence modulo two. *Mathematics of Computation.* 19(90): 201–209.

United Nations Officer for the Coordination of Humanitarian Affairs. Accessed March 30, 2016, http://www.unocha.org/syria.

United States Air Force (2003) *Air Force Doctrine Document (AFDD) 1* (Department of the Air Force, Washington, DC).

United States Department of the Army (2008) *The U.S. Army/Marine Corps Counterinsurgency Field Manual*. (Headquarters, Department of the Army, Washington, DC).

United States Joint Forces Command (2010) *Joint Operating Environment* (United States Joint Forces Command Center for Joint Futures, Suffolk, VA).

Van Rossum, Guido (2007) Python programming language. *USENIX Annual Technical Conference*, 41.

von Stackelberg, H (1952). *Marketform and Gleichgewicht* (Springer, Vienna). An English translation appeared in 1952 entitled *The Theory of the Market Economy* (Oxford University Press, Oxford).

Washburn A, Wood K (1995) Two-person zero-sum games for network interdiction. *Operations Research.* 43(2): 243–251.

Wollmer, R (1963) Some methods for determining the most vital link in a railway network. RAND Memorandum, RM-3321-ISA. April.

Wollmer R (1964) Removing arcs from a network. *Operations Research.* 12(6): 934–940.

Wood R (2011) Bilevel network interdiction models: Formulations and solutions. *Wiley Encyclopedia of Operations Research and Management Science*. (Wiley, Hoboken, NJ).

Zinkevich, M (2003) Online convex programming and generalized infinitesimal gradient ascent. *Proceedings of the Twentieth International Conference on Machine Learning* (AAAI Press, Menlo Park, CA) 928–936.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California